Domain Name System Operations
Internet-Draft
Updates: 10341035 (if approved)
Intended status: Standards Track
Expires: 18 July 2026

J. Abley
S. Neuteboom
Cloudflare

14 January 2026

Ordering of RRSets in DNS Message Sections
draft-jabley-dnsop-ordered-answer-section-00

Abstract

The existing Domain Name System (DNS) specifications lack some clarity in their description of the process by which individual sections of a DNS message are constructed.

This document updates RFC 1034 and RFC 1035 to provide a clearer specification, consistent with deployed implementations.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://ableyjoe.github.io/draft-jabley-dnsop-ordered-answer-section/draft-jabley-dnsop-ordered-answer-section.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-jabley-dnsop-ordered-answer-section/.

Discussion of this document takes place on the Domain Name System Operations Working Group mailing list (mailto:dnsop@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/dnsop/. Subscribe at https://www.ietf.org/mailman/listinfo/dnsop/.

Source for this draft and an issue tracker can be found at https://github.com/ableyjoe/draft-jabley-dnsop-ordered-answer-section.

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/.

Copyright Notice

Table of Contents

1. Introduction

[RFC1034] specifies an algorithm to follow when constructing a response to a DNS QUERY. This algorithm in some cases can result in multiple RRSets being included in a single section of a DNS message, e.g. when handling CNAME resource records.

Most consumers of DNS responses, such as stub resolvers, have interpreted the direction to copy or store particular RRSets in sections of a DNS response to mean "append", treating each section as an ordered list of RRSets. In particular, many stub stub resolvers are known to rely upon that interpretation when processing DNS responses, e.g. see Appendix A.

Some DNS implementations employ algorithms in other sections that aim to optimise processing of responses received by initiators, e.g. NAPTR before SRV before A/AAAA in the additional section of a response. This behaviour has not been observed to cause any interoperability problems, and is explicitly permitted by this document.

This document updates [RFC1035] to specify that the answer section in a DNS message is an ordered list of RRSets, but that other sections may be ordered differently. This document clarifies the directions provided in [RFC1034] to match the observed behaviour and expectations of deployed software.

2. Conventions and Definitions

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.

This document assumes familiarity with terminology specific to the Domain Name System (DNS) as described in [RFC9499].

3. Updates to RFC 1034

[RFC1034] specifies the algorithms by which sections of a DNS response are constructed. For example, step 3 of the algorithm described in [RFC1034] section 4.3.2 contains the direction "copy all RRs which match QTYPE into answer section".

In this case, and in all other cases where [RFC1034] specifies that particular RRSets be included in the answer section of a DNS message, the section MUST be treated as an ordered list of RRSets. When it is necessary to include new RRSets in a section of a DNS message that is under construction, those RRSets MUST be appended. The receiver of a DNS message MAY refuse to process DNS messages that have been constructed differently.

When constructing other sections of a DNS message, each section MAY be treated as a non-ordered list. A receiver of a DNS message MUST NOT reject a DNS message on the basis of the order of RRSets in those sections.

4. Updates to RFC 1035

In a DNS message, the answer section MUST be considered to be an ordered set of RRSets. All other sections in a DNS message MUST be considered to be a non-ordered set.

DNS implementations MUST construct each section in a DNS response according to the algorithms specified in [RFC1034], as clarified in Section 3.

5. Security Considerations

The recommendations contained in this document have no known security implications.

6. IANA Considerations

This document has no IANA actions.

7. References

7.1. 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/rfc/rfc1034>.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/rfc/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/rfc/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/rfc/rfc8174>.

[RFC9499] Hoffman, P. and K. Fujiwara, "DNS Terminology", BCP 219, RFC 9499, DOI 10.17487/RFC9499, March 2024, <https://www.rfc-editor.org/rfc/rfc9499>.

7.2. Informative References

[Cisco2026]
Cisco, "Cisco Business Switches Reboot with Fatal Error from DNSC Process", 8 January 2026, <https://www.cisco.com/c/en/us/support/docs/smb/switches/Catalyst-switches/kmgmt3846-cbs-reboot-with-fatal-error-from-dnsc-process.html>.

[Neuteboom2026]
Neuteboom, S., "What came first: the CNAME or the A record?", 14 January 2026, <https://blog.cloudflare.com/cname-a-record-order-dns-standards/>.

Appendix A. Events of 8 January 2026

Cloudflare operates a well-known public DNS resolver known as 1.1.1.1, after one of the IPv4 addresses associated with the service. On 8 January a software change in the 1.1.1.1 service had the unintential side-effect of changing the order in which RRSets were encoded in the answer section of DNS responses, in the case where constructing the responses involved CNAME processing. The previous ordering was as clarified in Section 3 and Section 4. The change in behaviour was not detected by a corresponding failure in a regression test, since the ordering in the answer section was not considered to be significant.

Following the software release, Cloudflare became aware of significant numbers of deployed DNS client implementations that were suffering from failure. In particular, the getanswer_r() function invoked by the getaddrinfo() function in glibc was found to fail to function, and some deployed ethernet switches were observed to reboot when trying to resolve the names of configured NTP servers [Cisco2026].

The impact associated with this event was particularly widespread because of the widespread use of the 1.1.1.1 resolver. However, the two examples of client implementations are also widely deployed in systems that may well be upgraded only infrequently (or never upgraded at all).

See [Neuteboom2026] for additional information.

Appendix B. Editorial Notes (remove before publication)

B.1. draft-jabley-dnsop-ordered-sections-00

Initial draft circulated for comment in 2015; subsequently expired.

B.2. draft-jabley-dnsop-ordered-answer-section-00

Draft revitalised following some operational excitement.

Added competent co-author.

Acknowledgments

The contributions of Mark Andrews and Paul Vixie to the original revision of this document are acknowledged.

Authors' Addresses

Joe Abley
Cloudflare
Email: jabley@cloudflare.com

Sebastiaan Neuteboom
Cloudflare
Email: sebastiaan@cloudflare.com