

Extensions to DNS (EDNS1)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

This document specifies a number of extensions within the Extended DNS framework defined by [EDNS0], including several new extended label types and the ability to ask multiple questions in a single request.

1 - Rationale and Scope

1.1. EDNS (see [RFC2671]) specifies an extension mechanism to DNS (see [RFC1035]) which provides for larger message sizes, additional label types, and new message flags.


1.2. This document makes use of the EDNS extension mechanisms to add the ability to ask multiple questions in a single request.

2 - Affected Protocol Elements

2.1. Multiple queries in a question section have not been supported in DNS due the applicability of some DNS Message Header flags (such as AA) and of the RCODE field only to a single QNAME, QTYPE, and QCLASS. Multiple questions per request are desirable, and some way of asking them must be made available.

3 - OPT pseudo-RR Flags and Options

3.1. The extended RCODE and flags are structured as follows:



VERSION Indicates the implementation level of whoever sets it. Full conformance with the draft standard version of this specification is version ``1.'' Note that both requestors and responders should set this to the highest level they implement, that responders should send back RCODE=BADVERS and that requestors should be prepared to probe using lower version numbers if they receive an RCODE=BADVERS.

FM ``First match'' flag. Notable only when multiple questions are present. If set in a request, questions will be processed in wire order and the first question whose answer would have NOERROR AND ANCOUNT>0 is treated as if it were the only question in the query message. Otherwise, questions can be processed in any order and all possible answer records will be included in the response. Response FM should be ignored by requestors.

RRD ``Recursion really desired'' flag. Notable only when a request is processed by an intermediate name server (``forwarder'') who is not authoritative for the zone containing QNAME, and where QTYPE=ANY or QDCOUNT>1. If set in a request, the intermediate name server can only answer using unexpired cached answers (either positive or negative) which were atomically acquired using (a) the same QTYPE or set of QTYPES present in the current question and whose TTLs were each minimized to the

smallest among them when first cached, and (b) the same FM and LM settings present in the current question.

Z Set to zero by senders and ignored by receivers, unless modified in a subsequent specification.

4 - Multiple Questions for QUERY

4.1. If QDCOUNT>1, multiple questions are present. All questions must be for the same QNAME and QCLASS; only the QTYPE is allowed to vary. It is an error for QDCOUNT>1 and any QTYPE=ANY or QCLASS=ANY.

4.2. RCODE and AA apply to all RRs in the answer section having the QNAME that is shared by all questions in the question section. AA applies to all matching answers, and will not be set unless the exact original request was processed by an authoritative server and the response forwarded in its entirety.

4.3. If a multiple question request is processed by an intermediate server and the authority server does not support multiple questions, the intermediate server must generate an answer iteratively by making multiple requests of the authority server. In this case, AA must never be set in the final answer due to lack of atomicity of the contributing authoritative responses.

4.4. If iteratively processing a multiple question request using an authority server which can only process single question requests, if any contributing request generates a SERVFAIL response, then the final response's RCODE should be SERVFAIL.

4.5. An authority server processing a query for which QDCOUNT>1 will respond with a delegation or referral if any of the multiple QTYPES present would yield such a response when QDCOUNT==1.

4.6. An initiator can infer the absence of any RRs for one of the QTYPES where QDCOUNT>1 if the response contains no RRs of that type but some RRs for one of the other QTYPES present.

[5](#) - References

- [RFC1035] P. Mockapetris, ``Domain Names - Implementation and Specification,`` [RFC 1035](#), USC/Information Sciences Institute, November 1987.
- [RFC2671] P. Vixie, ``Extension mechanisms for DNS (EDNS0),`` [RFC 2671](#), IETF DNSIND, September 1998

[6](#) - Author's Address

Paul Vixie
Internet Software Consortium
950 Charter Street
Redwood City, CA 94063
+1.650.779.7001
<vixie@isc.org>