

Network Working Group
Internet Draft
Expiration Date: August 1996

Robert Elz
University of Melbourne

Randy Bush
RGnet, Inc.

February 1996

Clarifications to the DNS Specification

[draft-ietf-dnsind-clarify-00.txt](#)

[1.](#) Status of this Memo

This document is an Internet-Draft. 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."

To learn the current status of any Internet-Draft, please check the "lid-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

[2.](#) Abstract

This draft considers some areas that have been identified as problems with the specification of the Domain Name System, and proposes remedies for the defects identified. Two separate issues are considered, IP packet header address usage from multi-homed servers, and TTLs in sets of records with the same name, class, and type.

Internet Draft

[draft-ietf-dnsind-clarify-00.txt](#)

February 1996

3. Introduction

Several problem areas in the Domain Name System specification have been noted through the years. This draft addresses two of them. The two issues here are independent. Those issues are the question of which source address a multi-homed DNS server should use when replying to a query, and the issue of differing TTLs for DNS records with the same label, class and type.

Suggestions for clarifications to the DNS specification to avoid the problems caused are made in this draft. The solutions proposed herein are intended to stimulate discussion. It is entirely possible that the sense of either may be reversed before the next iteration of this draft.

4. Server reply source address selection

Many DNS clients, in fact, most DNS clients, if not all, whether a server acting as a client for the purposes of recursive query resolution, or a resolver, expect that the address from which a reply is received via UDP will be the same address as that to which the query eliciting the reply was sent. This, along with the identifier (ID) in the reply is used for disambiguating replies, and filtering spurious responses. This may, or may not, have been intended when the DNS was designed, but is now a fact of life.

Some multi-homed hosts running DNS servers fail to anticipate this usage, and consequently send replies from the "wrong" source address, causing the reply to be discarded by the client.

To avoid these problems, servers when responding to queries using UDP must cause the reply to be sent with the source address field in the IP header set to the address that was in the destination address field of the IP header of the packet containing the query causing the response. If this would cause the response to be sent from an illegal IP address for sources, then the response must not be sent.

[Aside: An alternative would be to finish the previous sentence with "... may be sent from any legal IP address allocated to the server."]

5. Multiple TTLs in a Resource Record Set

DNS Resource Records (RRs) each have a label, class, type, and data. While it is meaningless for two records to ever have label, class, type and data all equal (servers should suppress such duplicates if encountered), it is possible for many record types to exist with the same label class and type, but with different data. Such a group of records is hereby defined to be a Resource Record Set (RRSet).

In all cases, a query for a specific (or non-specific) label, class, and type, will always return all records in the associated RRSet - whether that be one or more RRs, or the response shall be marked as "truncated" if the entire RRSet will not fit in the response.

Resource Records also each have a time to live (TTL). It is possible for the RRs in a RRSet to have different TTLs, however this has no known useful purpose, and can cause partial replies (not marked "truncated") from a caching server, where the TTLs for some of the RRs in the RRSet have expired, but not all have.

Consequently the use of differing TTLs in a RRSet is hereby deprecated, all TTLs in a RRSet should be the same.

Should a client receive a response containing RRs from an RRSet with TTLs not all equal, it should treat the RRs for all purposes as if all TTLs in the RRSet had been set to the value of the lowest TTL in the RRSet.

Servers never merge RRs from a response with RRs in their cache to form a RRSet, they must either ignore the RRs in the response, or use those to replace existing RRs from the cache, as appropriate. Consequently the issue of TTLs varying between the cache and a response does not cause concern, one will be ignored.

A Resource Record Set should only be included once in any DNS reply. It may occur in any of the Answer, Authority, or Additional

Information sections, as required, however should not be repeated in the same, or any other, section, except where explicitly required by a specification. Eg: an AXFR response requires the SOA record (always an RRSet containing a single RR) be both the first and last record of the reply. Where duplicates are required this way, the TTL transmitted in each case must be the same.

[6.](#) Security Considerations

This document does not consider security.

In particular, nothing in [section 4](#) is any way related to, or useful for, any security related purposes.

It is not believed that anything in this document adds to any security issues that may exist with the DNS, nor does it do anything to lessen them.

[7.](#) References

- [RFC1034] Domain Names - Concepts and Facilities,
P. Mockapetris, ISI, November 1987.
- [RFC1035] Domain Names - Implementation and Specification
P. Mockapetris, ISI, November 1987

[8.](#) Acknowledgements

To be supplied.

