

# RFC1035 - Domain Names - Implementation and Specification

[return to the IETF Standards](#)

Table 1: Data sheet for RFC1935 Domain Names - Implementation and Specification

|                 |                                                                                       |
|-----------------|---------------------------------------------------------------------------------------|
| Title           | Domain Names - Implementation and Specification                                       |
| Acronym         |                                                                                       |
| Version         | 1987                                                                                  |
| Document Number | RFC1035                                                                               |
| Release Date    | November 1987                                                                         |
| Reference       | <a href="https://tools.ietf.org/html/rfc1035">https://tools.ietf.org/html/rfc1035</a> |

**Note:** The following is an excerpt from the official IETF RFC. It is provided here as a convenience and is not authoritative. Refer to the original document as the authoritative reference.

## Introduction

*The goal of domain names is to provide a mechanism for naming resources in such a way that the names are usable in different hosts, networks, protocol families, internets, and administrative organizations.*

*From the user's point of view, domain names are useful as arguments to a local agent, called a resolver, which retrieves information associated with the domain name. Thus a user might ask for the host address or mail information associated with a particular domain name. To enable the user to request a particular type of information, an appropriate query type is passed to the resolver with the domain name. To the user, the domain tree is a single information space; the resolver is responsible for hiding the distribution of data among name servers from the user.*

*From the resolver's point of view, the database that makes up the domain space is distributed among various name servers. Different parts of the domain space are stored in different name servers, although a particular data item will be stored redundantly in two or more name servers. The resolver starts with knowledge of at least one name server. When the resolver processes a user query it asks a known name server for the information; in return, the resolver either receives the desired information or a referral to another name server. Using these referrals, resolvers learn the identities and contents of other name servers. Resolvers are responsible for dealing with the distribution of the domain space and dealing with the effects of name server failure by consulting redundant databases in other servers.*

*Name servers manage two kinds of data. The first kind of data held in sets called zones; each zone is the complete database for a particular "pruned" subtree of the domain space. This data is called authoritative. A name server periodically checks to make sure that its zones are up to date, and if not, obtains a new copy of updated zones from master files stored locally or in another name server. The second kind of data is cached data which was acquired by a local resolver. This data*

*may be incomplete, but improves the performance of the retrieval process when non-local data is repeatedly accessed. Cached data is eventually discarded by a timeout mechanism.*

*This functional structure isolates the problems of user interface, failure recovery, and distribution in the resolvers and isolates the database update and refresh problems in the name servers.*

From:  
<https://www.omgwiki.org/dido/> - **DIDO Wiki**

Permanent link:  
[https://www.omgwiki.org/dido/doku.php?id=dido:public:ra:xapend:xapend.b\\_stds:tech:ietf:1035&rev=1625163176](https://www.omgwiki.org/dido/doku.php?id=dido:public:ra:xapend:xapend.b_stds:tech:ietf:1035&rev=1625163176)

Last update: **2021/07/01 14:12**

