DNS for Clients

The Domain Name System (DNS) is a distributed database that's main function is to map IP addresses to host/domain names. The DNS system is made up of two types of users, servers and resolvers. Servers perform much of the work, storing local database information, answering questions from resolvers and servers, fetching information from other servers, and caching information for later use. Resolvers on the other hand, only talk to local servers that they are told to use, and are limited to two questions they can ask.

In this paper we are only going to look at how to build a resolver. We will assume functional name servers at our disposal, and a basic knowledge of what domain names, IP addresses, and IP/UDP packets are. We are also going to assume the server handles "recursion desired" requests (so the servers do all the work.)

The DNS Message format for both queries and responses is show on pg191 of Stevens vol 1.

The resolvers request message is very simple, it consists of a identification word that is simply a counter, a flags word 01 00 for a recursion desired query.

To resolve 'www.svn.net' we would send the following message to a name server @ port 53 :

```
00 00-01 00-00 01-00 00-00 00-00 00-03 77 77 77-03 73 76 6e-03 6e
65 74-00-00 01-00 01
```

Where 00 00 = identifier, usually a each message from a client will increment this value by one (the next message would use 00 01.

01 00 = flags field, we should always use 01 00 for a query, this asks the name server to do all the work (i.e. Recursion desired.)
00 01 = number of questions, we will always ask only 1 question.
00 00 = number of answer RR's, always 0.
00 00 = number of authority RR's, always 0.
00 00 = number of additional RR's, always 0.

--- the next info is the question portion of the DNS query message—
03 77 77 77 = www (length 3 + the ASCII chars www)
03 73 76 6e = svn
03 6e 65 74 = net
0 = no more name, terminates www.svn.net
00 01 = Query type, we want an IP address, use 00 05 if you want to resolve an IP address to a name.
00 01 = Query Class, always use 00 01 for IP.

Reply from nameserver:

```
0000  00 00-85 80-00 01-00 02-00 01-00 01-03 77 77 77-
0010  03 73 76 6e-03 6e 65 74-00-00 01-00 01-c0 0c-00
0020  05-00 01-00 01 51 80-00 0e-04 77 77 77 31-03 73
0030  76 6e-03 6e 65 74-00-c0 29-00 01-00 01-00 01 51
0040  00-00 04-cf 30 54 0f-c0 2e-00 02-00 01-00 01 51
0050  80-00 06-03 6e 73 31-c0 2e-00 02-00 01-00 01 51
0060  80-00 06-03 6e 73 31-c0 2e-c0 53-00 01-00 01-00
0070  01 51 80-00 04-cf 30 54 0a
```

On return we need to check the flags field, if the recursion available bit is not on, our server is not compatible with our resolver. If the recursion available bit is on then we want to check the last 8 bits of the flag for the response code, 0 means we have a valid response, 3 means there was an error resolving the name.

If we have no error then the address we are looking for is available in the data.

The above packet breaks down as follows:
00 00 = Identifier, matches the one we sent in the request.
85 80 = Flags, 85=response, query, authoritative answer, not tuncated.
= 80= unicast packet, recursion available, no error.
- we should only handle packet that are not truncated with recursion available and no error, else we should return a failed lookup.

00 01 = question count 1
00 02 = answer count 2
00 01 = authority count 1
00 01 = additional records 1

-Question section-
03 77 77 77 03 73 76 63 03 6e 65 74 00 =www.svn.net
00 01 = host address
00 01 = internet

-Answer section-
-answer 1-
c0 0c = www.svn.net (compressed) c0 means look at offset 0c which is the above 03 77 77 77… in the question section.
00 05 = canonical name for alias, (www.svn.net is an alias for www1.svn.net, see below)
00 01 = class internet
00 01 51 80 = time to live.
00 0e = lenth of objet that follows (length of below)
04 77 77 77 31 03 73 76 6e 03 6e 65 74 00 = www1.svn.net

-answer 2-
c0 29 = www1.svn.net (compressed)
00 01 = host address
00 01 = class internet
00 01 51 80 = time to live
00 04 = length of next object (following IP address)
cf 30 54 0f = IP address [207.48.84.15]

-Authority section-
c0 2e = svn.net (compressed)
00 02 = authoritative name server
00 01 = class internet
00 01 51 80 = time to live
00 06 = length of next object (below)
03 6e 73 31 c0 2e = ns1.svn.net (partually compressed)

-Additional record section 1 -
c0 53 = ns1.svn.net (compressed)
00 01 = host address
00 01=class internet
00 01 51 80= time to live
00 04 = length of following object (below IP)
cf 30 54 0a= IP address 207.48.84.10

All we really care about in this example is the 1st IP address for www1.svn.net [207.48.84.15], the rest is not relevant to our pursuits.

The complete record is shown only in that it is typical what to expect, the record could be longer if there are more namservers listed in the authority records.

To break it down in the simplest terms, ask one question, take the reply and check the flags, if the flags are what we want, find the first IP address record and return that. In a simpler case there might just be one answer (no alias), but we have to be able to deal with an alias.