(domain information groper) is a flexible command line tool which can be used to gather information from the Domain Name System servers. has two modes: simple interactive mode for a single query, and batch mode which executes a query for each in a list of several query lines. All query options are accessible from the command line. The usual simple use of will take the form: where: may be either a domain name or a dot-notation Internet address. If this optional field is omitted, will attempt to use the default name server for your machine.

If a domain name is specified, this will be resolved using the domain name system resolver (i.e., BIND). If your system does not support DNS, you may to specify a dot-notation address. Alternatively, if there is a server at your disposal somewhere, all that is required is that be present and indicate where the default name servers reside, so that itself can be resolved. See for information on Changing will affect both the standard resolver library and several programs which use it. As an option, the user may set the environment variable to name a file which is to be used instead of is specific to the resolver and is not referenced by the standard resolver If the variable is not set or the specified file is not readable, then will be used. is the domain name for which you are requesting information. See the option (documented in the subsection of this section) for convenient way to specify inverse address query. is the type of information (DNS query type) that you are requesting. If omitted, the default is The following types are recognized: network address all/any information about specified domain mail exchanger for the domain name servers zone of authority record host information zone transfer (must ask an authoritative server) arbitrary number of strings (See RFC 1035 for the complete list.) is the network class requested in the query. If omitted, the default is The following classes are recognized: Internet class domain all/any class information (See RFC 1035 for the complete list.) can be used to specify a and/or a of query. will parse the first occurrence of to mean To specify you must either specify twice, or set using the option (see below). is used to included an argument that is simply not parsed. This may be useful if running in batch mode. Instead of resolving every in a list of queries, you can avoid the overhead of doing so, and still have the domain name on the command line as a reference. Example: is used to specify an option which affects the operation of The following options are currently available (although not guaranteed to be useful): Convenient form to specify inverse address mapping. Instead of one can simply File for batch mode. The file contains a list of query specifications (command lines) which are to be executed successively. Lines beginning with or are ignored. Other options may still appear on command line, and will be in effect for each batch query. Time in seconds between start of successive queries when running in batch mode. Can be used to keep two or more batch commands running roughly in sync. Default is zero. Port number. Query a name server listening to a non-standard port number. Default is 53. After query returns, execute a command for response time comparison. This rather unelegantly makes a call to the shell. The last three lines of statistics is printed for the command: If the optional is present, it replaces in the shell command. Specify type of query. May specify either an integer value to be included in the type field or use the abbreviated mnemonic as discussed above (i.e., Specify class of query. May specify either an integer value to be included in the class field or use the abbreviated mnemonic as discussed above (i.e., in = C_IN). Sign the query with the TSIG key named keyname that is in the directory keydir. This flag specifies that the environment (defaults, print options, etc.), after all of the arguments are parsed, should be saved to a file to become the default environment. This is useful if you do not like the standard set of defaults and do not desire to include a large number of options each time is used. The environment consists of resolver state variable flags, timeout, and retries as well as the flags detailing output (see below). If the shell environment variable is set to the name of a file, this is where the default environment is saved. If not, the file is created in the current working directory. is specific to the resolver, and will not affect operation of the standard resolver library. Each time is executed, it looks for or the file specified by the shell environment variable If such file exists and is readable, then the environment is restored from this file before any arguments are parsed. This flag only affects batch query runs. When is specified on a line in a batch file, the environment after the arguments are parsed becomes the default environment for the duration of the batch file, or until the next line which specifies This flag only affects batch query runs. It specifies that the environment (as read initially or set by switch) is to be restored before each query (line) in a batch file. The default means that the environment does not stick, hence options specified on a single line in a batch file will remain in effect for subsequent lines (i.e. they are not restored to the default). is used to specify an option to be changed in the query packet or to change output specifics. Many of these are the same parameters accepted by If an option requires a parameter, the form is as follows: Most keywords can be abbreviated. Parsing of the options is very simplistic — a value must not be separated from its keyword by white space. The following keywords are currently available: Keyword Abbrev. Meaning [default] turn on/off debugging mode turn on/off extra debugging mode use/don't use recursive lookup set number of retries to # set timeout length to # seconds keep open option (implies vc) use/don't use virtual circuit use/don't use default domain name use/don't use domain search list set default domain name to ignore/don't ignore trunc. errors use/don't use primary server authoritative query only flag echo parsed arguments print query statistics print basic header print header flags print TTLs print class info print outgoing query print reply print question section print answer section print authoritative section print additional section set to default print flags set to minimal default print flags set print flags to # (# can be hex/octal/decimal) bitwise and print flags with # bitwise or print flags with # The and options affect the retransmission strategy used by the resolver library when sending datagram queries. The algorithm is as follows: for i = 0 to retry - 1

for j = 1 to num_servers send_query wait((time * (2**i)) / num_servers)

end end (Note: always uses a value of 1 for once required a slightly modified version of the BIND library. As of BIND 4.9, BIND's resolver has been augmented to work properly with Essentially, is a straight-forward (albeit not pretty) effort of parsing arguments and setting appropriate parameters. uses routines as well as accessing the structure. file to use in place of Pa /etc/resolv.conf default environment file See also the explanation of the and options, above. initial domain name and name server addresses default save file for default options RFC 1035. Steve Hotz

hotz@isi.edu uses functions from authored by Andrew Cherenson. has a serious case of "creeping featurism" -- the result of considering several potential uses during it's development. It would probably benefit from a rigorous diet. Similarly, the print flags and granularity of the items they specify make evident their rather ad hoc genesis. does not consistently exit nicely (with appropriate status) when a problem occurs somewhere in the resolver most of the common exit cases are handled This is particularly annoying when running in batch mode. If it exits abnormally (and is not caught), the entire batch aborts; when such an event is trapped, simply continues with the next query.