

Man Page - exim_dbmbuild - Build a DBM file.

- NAME

exim_dbmbuild - Build a DBM file.

- SYNOPSIS

```
exim_dbmbuild [-nolc] [-nozero] [-noduperr] [-nowarn] inputfile|- outputfile
```

- DESCRIPTION

The `exim_dbmbuild` program reads an input file containing keys and data in the format used by the `lsearch` lookup (see section 9.1). It writes a DBM file using the lower-cased alias names as keys and the remainder of the information as data. The lower-casing can be prevented by calling the program with the `-nolc` option.

A terminating zero is included as part of the key string. This is expected by the `dbm` lookup type. However, if the option `-nozero` is given, `exim_dbmbuild` creates files without terminating zeroes in either the key strings or the data strings. The `dbmz` lookup type can be used with such files.

The program requires two arguments: the name of the input file (which can be a single hyphen to indicate the standard input), and the name of the output file. It creates the output under a temporary name, and then renames it if all went well. If the native DB interface is in use (`USE_DB` is set in a compile-time configuration file - this is common in free versions of Unix) the two file names must be different, because in this mode the Berkeley DB functions create a single output file using exactly the name given. For example,

```
exim_dbmbuild /etc/aliases /etc/aliases.db
```

reads the system alias file and creates a DBM version of it in `/etc/aliases.db`.

In systems that use the `ndbm` routines (mostly proprietary versions of Unix), two files are used, with the suffixes `.dir` and `.pag`. In this environment, the suffixes are added to the second argument of `exim_dbmbuild`, so it can be the same as the first. This is also the case when the Berkeley functions are used in compatibility mode (though this is not recommended), because in that case it adds a `.db` suffix to the file name.

If a duplicate key is encountered, the program outputs a warning, and when it finishes, its return code is 1 rather than zero, unless the `-noduperr` option is used. By default, only the first of a set of duplicates is used - this makes it compatible with `lsearch` lookups. There is an option `-lastdup` which causes it to use the data for the last duplicate instead. There is also an option `-nowarn`, which stops it listing duplicate keys to "stderr". For other errors, where it doesn't actually make a new file, the return code is 2.