

# GNU dbm

A Database Manager

by Philip A. Nelson, Jason Downs and Sergey Poznyakoff

Manual by Pierre Gaumond, Philip A. Nelson, Jason Downs  
and Sergey Poznyakoff

Edition 1.11

for GNU dbm, Version 1.11

Published by the Free Software Foundation, 51 Franklin Street, Fifth Floor Boston, MA 02110-1301, USA

Copyright © 1989-1999, 2007-2011 Free Software Foundation, Inc.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover, and no Back-Cover texts. A copy of the license is included in the section entitled “GNU Free Documentation License.”

## Short Contents

1	Copying Conditions. . . . .	1
2	Introduction to GNU <code>dbm</code> . . . . .	2
3	List of functions. . . . .	3
4	Opening the database. . . . .	4
5	Closing the database. . . . .	5
6	Number of Records . . . . .	6
7	Inserting and replacing records in the database. . . . .	7
8	Searching for records in the database. . . . .	8
9	Removing records from the database. . . . .	9
10	Sequential access to records. . . . .	10
11	Database reorganization. . . . .	12
12	Database Synchronization . . . . .	13
13	Export and Import . . . . .	14
14	Error strings. . . . .	18
15	Setting options . . . . .	19
16	File Locking. . . . .	22
17	Useful global variables. . . . .	23
18	Error codes . . . . .	25
19	Compatibility with standard <code>dbm</code> and <code>ndbm</code> . . . . .	28
20	Examine and modify a GDBM database. . . . .	32
21	The <code>gdbm_dump</code> utility . . . . .	42
22	The <code>gdbm_load</code> utility . . . . .	43
23	Export a database into a portable format. . . . .	44
24	Exit codes . . . . .	45
25	Problems and bugs. . . . .	46
26	Additional resources . . . . .	47
A	GNU Free Documentation License. . . . .	48
	Index. . . . .	56

## Table of Contents

1	Copying Conditions.....	1
2	Introduction to GNU dbm.....	2
3	List of functions.....	3
4	Opening the database.....	4
5	Closing the database.....	5
6	Number of Records.....	6
7	Inserting and replacing records in the database.....	7
8	Searching for records in the database.....	8
9	Removing records from the database.....	9
10	Sequential access to records.....	10
11	Database reorganization.....	12
12	Database Synchronization.....	13
13	Export and Import.....	14
14	Error strings.....	18
15	Setting options.....	19
16	File Locking.....	22
17	Useful global variables.....	23
18	Error codes.....	25

<b>19</b>	<b>Compatibility with standard dbm and ndbm.</b>	
	.....	<b>28</b>
19.1	NDBM interface functions.....	28
19.2	DBM interface functions.....	30
<b>20</b>	<b>Examine and modify a GDBM database. . .</b>	<b>32</b>
20.1	gdbmtool invocation .....	32
20.2	gdbmtool interactive mode .....	33
20.2.1	Shell Variables .....	34
20.2.2	Gdbmtool Commands .....	36
20.2.3	Data Definitions .....	39
20.2.4	Startup Files .....	40
<b>21</b>	<b>The gdbm_dump utility .....</b>	<b>42</b>
<b>22</b>	<b>The gdbm_load utility .....</b>	<b>43</b>
<b>23</b>	<b>Export a database into a portable format.</b>	
	.....	<b>44</b>
<b>24</b>	<b>Exit codes .....</b>	<b>45</b>
<b>25</b>	<b>Problems and bugs .....</b>	<b>46</b>
<b>26</b>	<b>Additional resources .....</b>	<b>47</b>
<b>Appendix A GNU Free Documentation License</b>		
	.....	<b>48</b>
<b>Index .....</b>		<b>56</b>

# 1 Copying Conditions.

This library is *free*; this means that everyone is free to use it and free to redistribute it on a free basis. GNU `dbm` (`gdbm`) is not in the public domain; it is copyrighted and there are restrictions on its distribution, but these restrictions are designed to permit everything that a good cooperating citizen would want to do. What is not allowed is to try to prevent others from further sharing any version of `gdbm` that they might get from you.

Specifically, we want to make sure that you have the right to give away copies `gdbm`, that you receive source code or else can get it if you want it, that you can change these functions or use pieces of them in new free programs, and that you know you can do these things.

To make sure that everyone has such rights, we have to forbid you to deprive anyone else of these rights. For example, if you distribute copies `gdbm`, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must tell them their rights.

Also, for our own protection, we must make certain that everyone finds out that there is no warranty for anything in the `gdbm` distribution. If these functions are modified by someone else and passed on, we want their recipients to know that what they have is not what we distributed, so that any problems introduced by others will not reflect on our reputation.

`Gdbm` is currently distributed under the terms of the GNU General Public License, Version 3. (*NOT* under the GNU General Library Public License.) A copy the GNU General Public License is included with the distribution of `gdbm`.

## 2 Introduction to GNU dbm.

GNU dbm (**gdbm**) is a library of database functions that use extensible hashing and works similar to the standard UNIX **dbm** functions. These routines are provided to a programmer needing to create and manipulate a hashed database. (**gdbm** is *NOT* a complete database package for an end user.)

The basic use of **gdbm** is to store key/data pairs in a data file. Each key must be unique and each key is paired with only one data item. The keys can not be directly accessed in sorted order. The basic unit of data in **gdbm** is the structure:

```
typedef struct {
    char *dptr;
    int  dsize;
} datum;
```

This structure allows for arbitrary sized keys and data items.

The key/data pairs are stored in a **gdbm** disk file, called a **gdbm** database. An application must open a **gdbm** database to be able manipulate the keys and data contained in the database. **gdbm** allows an application to have multiple databases open at the same time. When an application opens a **gdbm** database, it is designated as a **reader** or a **writer**. A **gdbm** database can be opened by at most one writer at a time. However, many readers may open the database simultaneously. Readers and writers can not open the **gdbm** database at the same time.

### 3 List of functions.

The following is a quick list of the functions contained in the `gdbm` library. The include file `gdbm.h`, that can be included by the user, contains a definition of these functions.

```
#include <gdbm.h>

GDBM_FILE gdbm_open(name, block_size, flags, mode, fatal_func);
void gdbm_close(dbf);
int gdbm_store(dbf, key, content, flag);
datum gdbm_fetch(dbf, key);
int gdbm_delete(dbf, key);
datum gdbm_firstkey(dbf);
datum gdbm_nextkey(dbf, key);
int gdbm_reorganize(dbf);
void gdbm_sync(dbf);
int gdbm_exists(dbf, key);
char *gdbm_strerror(errno);
int gdbm_setopt(dbf, option, value, size);
int gdbm_fdesc(dbf);
int gdbm_export (GDBM_FILE, const char *, int, int);
int gdbm_export_to_file (GDBM_FILE dbf, FILE *fp);
int gdbm_import (GDBM_FILE, const char *, int);
int gdbm_import_from_file (GDBM_FILE dbf, FILE *fp, int flag);
int gdbm_count (GDBM_FILE dbf, gdbm_count_t *pcount);
int gdbm_version_cmp (int const a[], int const b[]);
```

The `gdbm.h` include file is often in the `/usr/local/include` directory. (The actual location of `gdbm.h` depends on your local installation of `gdbm`.)



## 4 Opening the database.

`GDBM_FILE gdbm_open (const char *name, int block_size, int flags, int mode, void (*fatal_func)(const char *))` [gdbm interface]

Initializes `gdbm` system. If the file has a size of zero bytes, a file initialization procedure is performed, setting up the initial structure in the file.

The arguments are:

*name*        The name of the file (the complete name, `gdbm` does not append any characters to this name).

*block\_size*   It is used during initialization to determine the size of various constructs. It is the size of a single transfer from disk to memory. This parameter is ignored if the file has been previously initialized. The minimum size is 512. If the value is less than 512, the file system block size is used, otherwise the value of *block\_size* is used.

*flags*        If *flags* is set to 'GDBM\_READER', the user wants to just read the database and any call to `gdbm_store` or `gdbm_delete` will fail. Many readers can access the database at the same time. If *flags* is set to 'GDBM\_WRITER', the user wants both read and write access to the database and requires exclusive access. If *flags* is set to 'GDBM\_WRCREAT', the user wants both read and write access to the database and wants it created if it does not already exist. If *flags* is set to 'GDBM\_NEWDB', the user want a new database created, regardless of whether one existed, and wants read and write access to the new database.

The following may also be logically or'd into the database flags: 'GDBM\_SYNC', which causes all database operations to be synchronized to the disk, 'GDBM\_NOLOCK', which prevents the library from performing any locking on the database file, and 'GDBM\_NOMMAP', which disables the memory mapping mechanism. The option 'GDBM\_FAST' is now obsolete, since `gdbm` defaults to no-sync mode.

If the host 'open' call (see [Section "open" in \*open\(2\) man page\*](#)) supports the 'O\_CLOEXEC' flag, the 'GDBM\_CLOEXEC' can be or'd into the flags, to enable the close-on-exec flag for the database file descriptor.

*mode*        File mode (see [Section "change permissions of a file" in \*chmod\(2\) man page\*](#), and see [Section "open a file" in \*open\(2\) man page\*](#)), which is used if the file is created).

*fatal\_func*   A function for `gdbm` to call if it detects a fatal error. The only parameter of this function is a string. If the value of 'NULL' is provided, `gdbm` will use a default function.

The return value, is the pointer needed by all other functions to access that `gdbm` file. If the return is the 'NULL' pointer, `gdbm_open` was not successful. The errors can be found in `gdbm_errno` variable (see [Chapter 17 \[Variables\]](#), [page 23](#)). Available error codes are discussed in [Chapter 18 \[Error codes\]](#), [page 25](#).

In all of the following calls, the parameter *dbf* refers to the pointer returned from `gdbm_open`.

## 5 Closing the database.

It is important that every file opened is also closed. This is needed to update the reader/writer count on the file:

`void gdbm_close (GDBM_FILE dbf)` [gdbm interface]

This function closes the `gdbm` file and frees all memory associated with it. The parameter is:

*dbf*            The pointer returned by `gdbm_open`.

## 6 Number of Records

`int gdbm_count (GDBM_FILE dbf, gdbm_count_t *pcount)` [gdbm interface]  
Counts number of records in the database *dbf*. On success, stores it in the memory location pointed to by *pcount* and return 0. On error, sets `gdbm_errno` (if relevant, also `errno`) and returns -1.

## 7 Inserting and replacing records in the database.

```
int gdbm_store (GDBM_FILE dbf, datum key, datum content,      [gdbm interface]
               int flag)
```

The function `gdbm_store` inserts or replaces records in the database.

The parameters are:

<i>dbf</i>	The pointer returned by <code>gdbm_open</code> .
<i>key</i>	The search key.
<i>content</i>	The data to be associated with the key.
<i>flag</i>	Defines the action to take when the key is already in the database. The value 'GDBM_REPLACE' (defined in <code>gdbm.h</code> ) asks that the old data be replaced by the new <i>content</i> . The value 'GDBM_INSERT' asks that an error be returned and no action taken if the <i>key</i> already exists.

This function can return the following values:

-1	The item was not stored in the database because the caller was not an official writer or either <i>key</i> or <i>content</i> have a 'NULL' 'dptr' field. Both <i>key</i> and <i>content</i> must have the 'dptr' field be a non-'NULL' value. Since a 'NULL' 'dptr' field is used by other functions to indicate an error, it cannot be valid data.
+1	The item was not stored because the argument <i>flag</i> was 'GDBM_INSERT' and the <i>key</i> was already in the database.
0	No error. The value of <i>content</i> is keyed by <i>key</i> . The file on disk is updated to reflect the structure of the new database before returning from this function.

If you store data for a *key* that is already in the data base, `gdbm` replaces the old data with the new data if called with 'GDBM\_REPLACE'. You do not get two data items for the same *key* and you do not get an error from `gdbm_store`.

The size in `gdbm` is not restricted like `dbm` or `ndbm`. Your data can be as large as you want.

## 8 Searching for records in the database.

**datum gdbm\_fetch** (*GDBM\_FILE dbf, datum key*) [gdbm interface]

Looks up a given *key* and returns the information associated with it. The ‘*dptr*’ field in the structure that is returned points to a memory block allocated by `malloc`. It is the caller’s responsibility to free it when no longer needed.

If the ‘*dptr*’ is ‘NULL’, no data was found.

The parameters are:

*dbf*            The pointer returned by `gdbm_open`.

*key*            The search key.

An example of using this function:

```
content = gdbm_fetch (dbf, key);
if (content.dptr == NULL)
{
    fprintf(stderr, "key not found\n");
}
else
{
    /* do something with content.dptr */
}
```

You may also search for a particular key without retrieving it:

**int gdbm\_exists** (*GDBM\_FILE dbf, datum key*) [gdbm interface]

Returns ‘true’ (‘1’) if the *key* exists in *dbf* and ‘false’ (‘0’) otherwise.

The parameters are:

*dbf*            The pointer returned by `gdbm_open`.

*key*            The search key.

## 9 Removing records from the database.

To remove some data from the database, use the `gdbm_delete` function.

`int gdbm_delete (GDBM_FILE dbf, datum key)` [gdbm interface]

Deletes the data associated with the given *key*, if it exists in the database *dbf*. The file on disk is updated to reflect the structure of the new database before returning from this function.

The parameters are:

*dbf*           The pointer returned by `gdbm_open`.

*datum key*    The search key.

The function returns ‘-1’ if the item is not present or the requester is a reader. The return of ‘0’ marks a successful delete.

## 10 Sequential access to records.

The next two functions allow for accessing all items in the database. This access is not key sequential, but it is guaranteed to visit every **key** in the database once. The order has to do with the hash values. `gdbm_firstkey` starts the visit of all keys in the database. `gdbm_nextkey` finds and reads the next entry in the hash structure for `dbf`.

**datum** `gdbm_firstkey` (*GDBM\_FILE dbf*) [gdbm interface]

Initiate sequential access to the database *dbf*. The returned value is the first key accessed in the database. If the ‘`dptr`’ field in the returned datum is ‘NULL’, the database contains no data.

Otherwise, ‘`dptr`’ points to a memory block obtained from `malloc`, which holds the key value. The caller is responsible for freeing this memory block when no longer needed.

**datum** `gdbm_nextkey` (*GDBM\_FILE dbf, datum prev*) [gdbm interface]

This function continues the iteration over the keys in *dbf*, initiated by `gdbm_firstkey`. The parameter *prev* holds the value returned from a previous call to `gdbm_nextkey` or `gdbm_firstkey`.

The function returns next key from the database. If the ‘`dptr`’ field in the returned datum is ‘NULL’, all keys in the database has been visited.

Otherwise, ‘`dptr`’ points to a memory block obtained from `malloc`, which holds the key value. The caller is responsible for freeing this memory block when no longer needed.

These functions were intended to visit the database in read-only algorithms, for instance, to validate the database or similar operations. The usual algorithm for sequential access is:

```
key = gdbm_firstkey (dbf);
while (key.dptr)
{
    datum nextkey;

    /* do something with the key */
    ...

    /* Obtain the next key */
    nextkey = gdbm_nextkey (dbf, key);
    /* Reclaim the memory used by the key */
    free (key.dptr);
    /* Use nextkey in the next iteration. */
    key = nextkey;
}
```

Care should be taken when the `gdbm_delete` function is used in such a loop. File visiting is based on a *hash table*. The `gdbm_delete` function re-arranges the hash table to make sure that any collisions in the table do not leave some item *un-findable*. The original key order is *not* guaranteed to remain unchanged in all instances. So it is possible that some key will not be visited if a loop like the following is executed:

```
key = gdbm_firstkey (dbf);
while (key.dptr)
{
    datum nextkey;
    if (some condition)
    {
        gdbm_delete (dbf, key);
    }
    nextkey = gdbm_nextkey (dbf, key);
    free (key.dptr);
    key = nextkey;
}
```



## 11 Database reorganization.

The following function should be used very seldom.

```
int gdbm_reorganize (GDBM_FILE dbf) [gdbm interface]
```

Reorganizes the database.

The parameter is:

*dbf*            The pointer returned by `gdbm_open`.

If you have had a lot of deletions and would like to shrink the space used by the `gdbm` file, this function will reorganize the database. This results, in particular, in shortening the length of a `gdbm` file by removing the space occupied by deleted records.

This reorganization requires creating a new file and inserting all the elements in the old file *dbf* into the new file. The new file is then renamed to the same name as the old file and *dbf* is updated to contain all the correct information about the new file. If an error is detected, the return value is negative. The value zero is returned after a successful reorganization.

## 12 Database Synchronization

Unless your database was opened with the ‘GDBM\_SYNC’ flag, `gdbm` does not wait for writes to be flushed to the disk before continuing. This allows for faster writing of databases at the risk of having a corrupted database if the application terminates in an abnormal fashion. The following function allows the programmer to make sure the disk version of the database has been completely updated with all changes to the current time.

`void gdbm_sync (GDBM_FILE dbf)` [gdbm interface]  
Synchronizes the changes in *dbf* with its disk file. The parameter is a pointer returned by `gdbm_open`.

This function would usually be called after a complete set of changes have been made to the database and before some long waiting time. The `gdbm_close` function automatically calls the equivalent of `gdbm_sync` so no call is needed if the database is to be closed immediately after the set of changes have been made.

## 13 Export and Import

Gdbm databases can be converted into so-called *flat format* files. Such files cannot be used for searching, their sole purpose is to keep the data from the database for restoring it when the need arrives. There are two flat file formats, which differ in the way they represent the data and in the amount of meta-information stored. Both formats can be used, for example, to migrate between the different versions of `gdbm` databases. Generally speaking, flat files are safe to send over the network, and can be used to recreate the database on another machine. The recreated database is guaranteed to be a byte-to-byte equivalent of the database from which the flat file was created. This does not necessarily mean, however, that this file can be used in the same way as the original one. For example, if the original database contained non-ASCII data (e.g. C structures, integers etc.), the recreated database can be of any use only if the target machine has the same integer size and byte ordering as the source one and if its C compiler uses the same packing conventions as the one which generated C which populated the original database. In general, such binary databases are not portable between machines, unless you follow some stringent rules on what data is written to them and how it is interpreted.

The GDBM version 1.11 supports two flat file formats. The *binary* flat file format was first implemented in GDBM version 1.9.1. This format stores only key/data pairs, it does not keep information about the database file itself. As its name implies, files in this format are binary files.

The *ascii* flat file format encodes all data in base64 and stores not only key/data pairs, but also the original database file metadata, such as file name, mode and ownership. Files in this format can be sent without additional encapsulation over transmission channels that normally allow only ASCII data, such as, e.g. SMTP. Due to additional metadata they allow for restoring an exact copy of the database, including file ownership and privileges, which is especially important if the database in question contained some security-related data.

We call a process of creating a flat file from a database *exporting* or *dumping* this database. The reverse process, creating the database from a flat file is called *importing* or *loading* the database.

```
int gdbm_dump (GDBM_FILE dbf, const char *filename, int          [gdbm interface]
               format, int open_flags, int mode)
```

Dumps the database file to the named file in requested format. Arguments are:

- |                   |  |
|-------------------|--|
| <i>dbf</i>        | A pointer to the source database, returned by a prior call to <code>gdbm_open</code> .   |
| <i>filename</i>   | Name of the dump file.   |
| <i>format</i>     | Output file format. Allowed values are: ‘GDBM_DUMP_FMT_BINARY’ to create a binary dump and ‘GDBM_DUMP_FMT_ASCII’ to create an ASCII dump file.   |
| <i>open_flags</i> | How to create the output file. If <i>flag</i> is ‘GDBM_WRCREAT’ the file will be created if it does not exist. If it does exist, the <code>gdbm_dump</code> will fail.<br>If <i>flag</i> is ‘GDBM_NEWDB’, the function will create a new output file, replacing it if it already exists. |
| <i>mode</i>       | The permissions to use when creating the output file. See <a href="#">Section “open a file” in <code>open(2)</code> man page</a> , for a detailed discussion.  |

```
int gdbm_load (GDBM_FILE *pdbf, const char *filename, int      [gdbm interface]
               flag, int meta_mask, unsigned long *errline)
```

Loads data from the dump file *filename* into the database pointed to by *pdbf*. The latter can point to 'NULL', in which case the function will try to create a new database. If it succeeds, the function will return, in the memory location pointed to by *pdbf*, a pointer to the newly created database. If the dump file carries no information about the original database file name, the function will set `gdbm_errno` to 'GDBM\_NO\_DBNAME' and return '-1', indicating failure.

The *flag* has the same meaning as the *flag* argument to the `gdbm_store` function (see [Chapter 7 \[Store\]](#), page 7).

The *meta\_mask* argument can be used to disable restoring certain bits of file's meta-data from the information in the input dump file. It is a binary OR of zero or more of the following:

GDBM\_META\_MASK\_MODE

Do not restore file mode.

GDBM\_META\_MASK\_OWNER

Do not restore file owner.

The function returns 0 upon successful completion or -1 on fatal errors and 1 on mild (non-fatal) errors.

If a fatal error occurs, `gdbm_errno` will be set to one of the following values:

GDBM\_FILE\_OPEN\_ERROR

Input file (*filename*) cannot be opened. The `errno` variable can be used to get more detail about the failure.

GDBM\_MALLOC\_ERROR

Not enough memory to load data.

GDBM\_FILE\_READ\_ERROR

Reading from *filename* failed. The `errno` variable can be used to get more detail about the failure.

GDBM\_ILLEGAL\_DATA

Input contained some illegal data.

GDBM\_ITEM\_NOT\_FOUND

This error can occur only when the input file is in ASCII format. It indicates that the data part of the record about to be read lacked length specification. Application developers are advised to treat this error equally as 'GDBM\_ILLEGAL\_DATA'.

Mild errors mean that the function was able to successfully load and restore the data, but was unable to change database file metadata afterward. The table below lists possible values for `gdbm_errno` in this case. To get more detail, inspect the system `errno` variable.

GDBM\_ERR\_FILE\_OWNER

The function was unable to restore database file owner.

## GDBM\_ERR\_FILE\_MODE

The function was unable to restore database file mode (permission bits).

If an error occurs while loading data from an input file in ASCII format, the number of line in which the error occurred will be stored in the location pointed to by the *errline* parameter, unless it is 'NULL'.

If the line information is not available or applicable, *errline* will be set to '0'.

```
int gdbm_dump_to_file (GDBM_FILE dbf, FILE *fp, int          [gdbm interface]
                      format)
```

This is an alternative entry point to `gdbm_dump` (which see). Arguments are:

*dbf*            A pointer to the source database, returned by a call to `gdbm_open`.

*fp*            File to write the data to.

*format*        Format of the dump file. See the *format* argument to the `gdbm_dump` function.

```
int gdbm_load_from_file (GDBM_FILE *pdbf, FILE *fp, int      [gdbm interface]
                        replace, int meta_mask, unsigned long *line)
```

This is an alternative entry point to `gdbm_dump`. It writes the output to *fp* which must be a file open for writing. The rest of arguments is the same as for `gdbm_load` (excepting of course *flag*, which is not needed in this case).

```
int gdbm_export (GDBM_FILE dbf, const char *exportfile,      [gdbm interface]
                int flag, int mode)
```

This function is retained for compatibility with GDBM 1.10 and earlier. It dumps the database to a file in binary dump format and is entirely equivalent to

```
gdbm_dump(dbf, exportfile, GDBM_DUMP_FMT_BINARY,
          flag, mode)
```

```
int gdbm_export_to_file (GDBM_FILE dbf, FILE *fp)           [gdbm interface]
```

This is an alternative entry point to `gdbm_export`. This function writes to file *fp* a binary dump of the database *dbf*.

```
int gdbm_import (GDBM_FILE dbf, const char *importfile,     [gdbm interface]
                int flag)
```

This function is retained for compatibility with GDBM 1.10 and earlier. It loads the file *importfile*, which must be a binary flat file, into the database *dbf* and is equivalent to the following construct:

```
dbf = gdbm_open (importfile, 0,
                flag == GDBM_REPLACE ?
                  GDBM_WRCREAT : GDBM_NEWDB,
                  0600, NULL);
gdbm_load (&dbf, exportfile, 0, flag, NULL)
```

```
int gdbm_import_from_file (GDBM_FILE dbf, FILE *fp, int      [gdbm interface]
                          flag)
```

An alternative entry point to `gdbm_import`. Reads the binary dump from the file *fp* and stores the key/value pairs to *dbf*. See [Chapter 7 \[Store\], page 7](#), for a description of *flag*.

This function is equivalent to:

```
dbf = gdbm_open (importfile, 0,  
                flag == GDBM_REPLACE ?  
                GDBM_WRCREAT : GDBM_NEWDB,  
                0600, NULL);  
gdbm_load_from_file (dbf, fp, flag, 0, NULL);
```

## 14 Error strings.

To convert a `gdbm` error code into English text, use this routine:

```
const char * gdbm_strerror (gdbm_error errno) [gdbm interface]  
    Converts errno (which is an integer value) into a human-readable descriptive text.  
    Returns a pointer to a static string. The caller must not alter or free the returned  
    pointer.
```

The *errno* argument is usually the value of the global variable `gdbm_errno`. See [Chapter 17 \[Variables\]](#), page 23.

## 15 Setting options

Gdbm supports the ability to set certain options on an already open database.

```
int gdbm_setopt (GDBM_FILE dbf, int option, void *value,      [gdbm interface]
                 int size)
```

Sets an option on the database or returns the value of an option.

The parameters are:

<i>dbf</i>	The pointer returned by <code>gdbm_open</code> .
<i>option</i>	The option to be set or retrieved.
<i>value</i>	A pointer to the value to which <i>option</i> will be set or where to place the option value (depending on the option).
<i>size</i>	The length of the data pointed to by <i>value</i> .

The valid options are:

GDBM.SETCACHESIZE

GDBM.CACHESIZE

Set the size of the internal bucket cache. This option may only be set once on each GDBM\_FILE descriptor, and is set automatically to 100 upon the first access to the database. The *value* should point to a `size_t` holding the desired cache size.

The ‘GDBM.CACHESIZE’ option is provided for compatibility with earlier versions.

GDBM.GETCACHESIZE

Return the size of the internal bucket cache. The *value* should point to a `size_t` variable, where the size will be stored.

GDBM.GETFLAGS

Return the flags describing the state of the database. The *value* should point to a `int` variable where to store the flags. The return is the same as the flags used when opening the database (see [Chapter 4 \[Open\], page 4](#)), except that it reflects the current state (which may have been altered by another calls to `gdbm_setopt`).

GDBM.FASTMODE

Enable or disable the *fast writes mode*, i.e. writes without subsequent synchronization. The *value* should point to an integer: ‘TRUE’ to enable fast mode, and ‘FALSE’ to disable it.

This option is retained for compatibility with previous versions of `gdbm`. Its effect is the reverse of `GDBM.SETSYNCMODE` (see below).

GDBM.SETSYNCMODE

GDBM.SYNCMODE

Turn on or off file system synchronization operations. This setting defaults to off. The *value* should point to an integer: ‘TRUE’ to turn synchronization on, and ‘FALSE’ to turn it off.



Note, that this option is a reverse of `GDBM_FASTMODE`, i.e. calling `GDBM_SETSYNCMODE` with ‘TRUE’ has the same effect as calling `GDBM_FASTMODE` with ‘FALSE’.

The ‘`GDBM_SYNCMODE`’ option is provided for compatibility with earlier versions.

#### `GDBM_GETSYNCMODE`

Return the current synchronization status. The *value* should point to an `int` where the status will be stored.

#### `GDBM_SETCENTFREE`

##### `GDBM_CENTFREE`

*NOTICE: This feature is still under study.*

Set central free block pool to either on or off. The default is off, which is how previous versions of `gdbm` handled free blocks. If set, this option causes all subsequent free blocks to be placed in the *global* pool, allowing (in theory) more file space to be reused more quickly. The *value* should point to an integer: ‘TRUE’ to turn central block pool on, and ‘FALSE’ to turn it off.

The ‘`GDBM_CENTFREE`’ option is provided for compatibility with earlier versions.

#### `GDBM_SETCOALESCEBLKS`

##### `GDBM_COALESCEBLKS`

*NOTICE: This feature is still under study.*

Set free block merging to either on or off. The default is off, which is how previous versions of `gdbm` handled free blocks. If set, this option causes adjacent free blocks to be merged. This can become a CPU expensive process with time, though, especially if used in conjunction with `GDBM_CENTFREE`. The *value* should point to an integer: ‘TRUE’ to turn free block merging on, and ‘FALSE’ to turn it off.

#### `GDBM_GETCOALESCEBLKS`

Return the current status of free block merging. The *value* should point to an `int` where the status will be stored.

#### `GDBM_SETMAXMAPSIZE`

Sets maximum size of a memory mapped region. The *value* should point to a value of type `size_t`, `unsigned long` or `unsigned`. The actual value is rounded to the nearest page boundary (the page size is obtained from `sysconf(_SC_PAGESIZE)`).

#### `GDBM_GETMAXMAPSIZE`

Return the maximum size of a memory mapped region. The *value* should point to a value of type `size_t` where to return the data.

#### `GDBM_SETTMAP`

Enable or disable memory mapping mode. The *value* should point to an integer: ‘TRUE’ to enable memory mapping or ‘FALSE’ to disable it.

#### `GDBM_GETTMAP`

Check whether memory mapping is enabled. The *value* should point to an integer where to return the status.

**GDBM.GETDBNAME**

Return the name of the database disk file. The *value* should point to a variable of type `char**`. A pointer to the newly allocated copy of the file name will be placed there. The caller is responsible for freeing this memory when no longer needed. For example:

```
char *name;

if (gdbm_setopt (dbf, GDBM_GETDBNAME, &name, sizeof (name)))
{
    fprintf (stderr, "gdbm_setopt failed: %s\n",
            gdbm_strerror (gdbm_errno));
}
else
{
    printf ("database name: %s\n", name);
    free (name);
}
```

The return value will be ‘-1’ upon failure, or ‘0’ upon success. The global variable `gdbm_errno` will be set upon failure.

For instance, to set a database to use a cache of 10, after opening it with `gdbm_open`, but prior to accessing it in any way, the following code could be used:

```
int value = 10;
ret = gdbm_setopt (dbf, GDBM_CACHESIZE, &value, sizeof (int));
```

## 16 File Locking.

With locking disabled (if `gdbm_open` was called with ‘`GDBM_NOLOCK`’), the user may want to perform their own file locking on the database file in order to prevent multiple writers operating on the same file simultaneously.

In order to support this, the `gdbm_fdesc` routine is provided.

`int gdbm_fdesc (GDBM_FILE dbf)` [gdbm interface]  
Returns the file descriptor of the database *dbf*. This value can be used as an argument to `flock`, `lockf` or similar calls.

## 17 Useful global variables.

The following global variables and constants are available:

`gdbm_error` `gdbm_errno` [Variable]

This variable contains error code from the last failed `gdbm` call. See [Chapter 18 \[Error codes\]](#), [page 25](#), for a list of available error codes and their descriptions.

Use `gdbm_strerror` (see [Chapter 14 \[Errors\]](#), [page 18](#)) to convert it to a descriptive text.

`const char * gdbm_errlist[]` [Variable]

This variable is an array of error descriptions, which is used by `gdbm_strerror` to convert error codes to human-readable text (see [Chapter 14 \[Errors\]](#), [page 18](#)). You can access it directly, if you wish so. It contains `_GDBM_MAX_ERRNO + 1` elements and can be directly indexed by the error code to obtain a corresponding descriptive text.

`_GDBM_MIN_ERRNO` [Constant]

The minimum error code used by `gdbm`.

`_GDBM_MAX_ERRNO` [Constant]

The maximum error code used by `gdbm`.

`const char * gdbm_version` [Variable]

A string containing the version information.

`int const gdbm_version_number[3]` [Variable]

This variable contains the `gdbm` version numbers:

Index	Meaning
0	Major number
1	Minor number
2	Patchlevel number

Additionally, the following constants are defined in the `gdbm.h` file:

`GDBM_VERSION_MAJOR`

Major number.

`GDBM_VERSION_MINOR`

Minor number.

`GDBM_VERSION_PATCH`

Patchlevel number.

These can be used to verify whether the header file matches the library.

To compare two split-out version numbers, use the following function:

`int gdbm_version_cmp (int const a[3], int const b[3])` [gdbm interface]

Compare two version numbers. Return ‘-1’ if *a* is less than *b*, ‘1’ if *a* is greater than *b* and ‘0’ if they are equal.

Comparison is done from left to right, so that:

```
a = { 1, 8, 3 };  
b = { 1, 8, 3 };  
gdbm_version_cmp (a, b)  $\Rightarrow$  0
```

```
a = { 1, 8, 3 };  
b = { 1, 8, 2 };  
gdbm_version_cmp (a, b)  $\Rightarrow$  1
```

```
a = { 1, 8, 3 };  
b = { 1, 9. 0 };  
gdbm_version_cmp (a, b)  $\Rightarrow$  -1
```

## 18 Error codes

This chapter summarizes error codes which can be set by the functions in `gdbm` library.

### `GDBM_NO_ERROR`

No error occurred.

### `GDBM_MALLOC_ERROR`

Memory allocation failed. Not enough memory.

### `GDBM_BLOCK_SIZE_ERROR`

This error is set by the `gdbm_open` function (see [Chapter 4 \[Open\], page 4](#)), if the value of its *block\_size* argument is incorrect.

### `GDBM_FILE_OPEN_ERROR`

The library was not able to open a disk file. This can be set by `gdbm_open` (see [Chapter 4 \[Open\], page 4](#)), `gdbm_export` and `gdbm_import` functions (see [Chapter 13 \[Flat files\], page 14](#)).

Inspect the value of the system `errno` variable to get more detailed diagnostics.

### `GDBM_FILE_WRITE_ERROR`

Writing to a disk file failed. This can be set by `gdbm_open` (see [Chapter 4 \[Open\], page 4](#)), `gdbm_export` and `gdbm_import` functions.

Inspect the value of the system `errno` variable to get more detailed diagnostics.

### `GDBM_FILE_SEEK_ERROR`

Positioning in a disk file failed. This can be set by `gdbm_open` (see [Chapter 4 \[Open\], page 4](#)) function.

Inspect the value of the system `errno` variable to get a more detailed diagnostics.

### `GDBM_FILE_READ_ERROR`

Reading from a disk file failed. This can be set by `gdbm_open` (see [Chapter 4 \[Open\], page 4](#)), `gdbm_export` and `gdbm_import` functions.

Inspect the value of the system `errno` variable to get a more detailed diagnostics.

### `GDBM_BAD_MAGIC_NUMBER`

The file given as argument to `gdbm_open` function is not a valid `gdbm` file: it has a wrong magic number.

### `GDBM_EMPTY_DATABASE`

The file given as argument to `gdbm_open` function is not a valid `gdbm` file: it has zero length.

### `GDBM_CANT_BE_READER`

This error code is set by the `gdbm_open` function if it is not able to lock file when called in ‘`GDBM_READER`’ mode (see [Chapter 4 \[Open\], page 4](#)).

### `GDBM_CANT_BE_WRITER`

This error code is set by the `gdbm_open` function if it is not able to lock file when called in writer mode (see [Chapter 4 \[Open\], page 4](#)).

**GDBM\_READER\_CANT\_DELETE**

Set by the `gdbm_delete` (see [Chapter 9 \[Delete\]](#), page 9) if it attempted to operate on a database that is open in read-only mode (see [Chapter 4 \[Open\]](#), page 4).

**GDBM\_READER\_CANT\_STORE**

Set by the `gdbm_store` (see [Chapter 7 \[Store\]](#), page 7) if it attempted to operate on a database that is open in read-only mode (see [Chapter 4 \[Open\]](#), page 4).

**GDBM\_READER\_CANT\_REORGANIZE**

Set by the `gdbm_reorganize` (see [Chapter 11 \[Reorganization\]](#), page 12) if it attempted to operate on a database that is open in read-only mode (see [Chapter 4 \[Open\]](#), page 4).

**GDBM\_UNKNOWN\_UPDATE**

Currently unused. Reserved for future uses.

**GDBM\_ITEM\_NOT\_FOUND**

Requested item was not found. This error is set by `gdbm_delete` (see [Chapter 9 \[Delete\]](#), page 9) and `gdbm_fetch` (see [Chapter 8 \[Fetch\]](#), page 8) when the requested key value is not found in the database.

**GDBM\_REORGANIZE\_FAILED**

The `gdbm_reorganize` function is not able to create a temporary database. See [Chapter 11 \[Reorganization\]](#), page 12.

**GDBM\_CANNOT\_REPLACE**

Cannot replace existing item. This error is set by the `gdbm_store` if the requested key value is found in the database and the *flag* parameter is not 'GDBM\_REPLACE'. See [Chapter 7 \[Store\]](#), page 7, for a detailed discussion.

**GDBM\_ILLEGAL\_DATA**

Either *key* or *content* parameter was wrong in a call to `gdbm_store` (see [Chapter 7 \[Store\]](#), page 7).

**GDBM\_OPT\_ALREADY\_SET**

Requested option can be set only once and was already set. This error is returned by the `gdbm_setopt` function. See [Chapter 15 \[Options\]](#), page 19.

**GDBM\_OPT\_ILLEGAL**

The *option* argument is not valid or the *value* argument points to an invalid value in a call to `gdbm_setopt` function. See [Chapter 15 \[Options\]](#), page 19.

**GDBM\_BYTE\_SWAPPED**

The `gdbm_open` function (see [Chapter 4 \[Open\]](#), page 4) attempts to open a database which is created on a machine with different byte ordering.

**GDBM\_BAD\_FILE\_OFFSET**

The `gdbm_open` function (see [Chapter 4 \[Open\]](#), page 4) sets this error code if the file it tries to open has a wrong magic number.

**GDBM\_BAD\_OPEN\_FLAGS**

Set by the `gdbm_export` function if supplied an invalid *flags* argument. See [Chapter 13 \[Flat files\]](#), page 14.

**GDBM\_FILE\_STAT\_ERROR**

Getting information about a disk file failed. The system `errno` will give more details about the error.

This error can be set by the following functions: `gdbm_open`, `gdbm_reorganize`.

**GDBM\_FILE\_EOF**

End of file was encountered where more data was expected to be present. This error can occur when fetching data from the database and usually means that the database is truncated or otherwise corrupted.

This error can be set by any GDBM function that does I/O. Some of these functions are: `gdbm_delete`, `gdbm_exists`, `gdbm_fetch`, `gdbm_export`, `gdbm_import`, `gdbm_reorganize`, `gdbm_firstkey`, `gdbm_nextkey`, `gdbm_store`.

**GDBM\_NO\_DBNAME**

Output database name is not specified. This error code is set by `gdbm_load` (see [\[gdbm\\_load\]](#), [page 15](#)) if the first argument points to 'NULL' and the input file does not specify the database name.

**GDBM\_ERR\_FILE\_OWNER**

This error code is set by `gdbm_load` if it is unable to restore database file owner. It is a mild error condition, meaning that the data have been restored successfully, only changing the target file owner failed. Inspect the system `errno` variable to get a more detailed diagnostics.

**GDBM\_ERR\_FILE\_MODE**

This error code is set by `gdbm_load` if it is unable to restore database file mode. It is a mild error condition, meaning that the data have been restored successfully, only changing the target file owner failed. Inspect the system `errno` variable to get a more detailed diagnostics.



## 19 Compatibility with standard `dbm` and `ndbm`.

`Gdbm` includes a compatibility layer, which provides traditional ‘`ndbm`’ and older ‘`dbm`’ functions. The layer is compiled and installed if the `--enable-libgdbm-compat` option is used when configuring the package.

The compatibility layer consists of two header files: `ndbm.h` and `dbm.h` and the `libgdbm_compat` library.

Older programs using `ndbm` or `dbm` interfaces can use `libgdbm_compat` without any changes. To link a program with the compatibility library, add the following two options to the `cc` invocation: `-lgdbm_compat -lgdbm`. The `-L` option may also be required, depending on where `gdbm` is installed, e.g.:

```
cc ... -L/usr/local/lib -lgdbm_compat -lgdbm
```

Please note that the compatibility library contains references to `gdbm` routines so the order in which the libraries are linked is essential. This means that the library linking order given in the above example must be respected.

Databases created and manipulated by the compatibility interfaces consist of two different files: `file.dir` and `file.pag`. This is required by the POSIX specification and corresponds to the traditional usage. Note, however, that despite the similarity of the naming convention, actual data stored in these files has not the same format as in the databases created by other `dbm` or `ndbm` libraries. In other words, you cannot access a standard UNIX `dbm` file with GNU `dbm`!

GNU `dbm` files are not `sparse`. You can copy them with the usual `cp` command and they will not expand in the copying process.

### 19.1 NDBM interface functions.

The functions below implement the POSIX ‘`ndbm`’ interface:

**DBM \* `dbm_open` (*char \*file*, *int flags*, *int mode*)** [ndbm]

Opens a database. The *file* argument is the full name of the database file to be opened. The function opens two files: *file.pag* and *file.dir*. The *flags* and *mode* arguments have the same meaning as the second and third arguments of `open` (see [Section “open a file” in \*open\(2\) man page\*](#)), except that a database opened for write-only access opens the files for read and write access and the behavior of the `O_APPEND` flag is unspecified.

The function returns a pointer to the DBM structure describing the database. This pointer is used to refer to this database in all operations described below.

Any error detected will cause a return value of ‘`NULL`’ and an appropriate value will be stored in `gdbm_errno` (see [Chapter 17 \[Variables\]](#), [page 23](#)).

**void `dbm_close` (*DBM \*dbf*)** [ndbm]

Closes the database. The *dbf* argument must be a pointer returned by an earlier call to `dbm_open`.

**datum `dbm_fetch` (*DBM \*dbf*, *datum key*)** [ndbm]

Reads a record from the database with the matching key. The *key* argument supplies the key that is being looked for.

If no matching record is found, the `dptr` member of the returned datum is `'NULL'`. Otherwise, the `dptr` member of the returned datum points to the memory managed by the compatibility library. The application should never free it.

`int dbm_store (DBM *dbf, datum key, datum content, int mode)` [ndbm]

Writes a key/value pair to the database. The argument `dbf` is a pointer to the DBM structure returned from a call to `dbm_open`. The `key` and `content` provide the values for the record key and content. The `mode` argument controls the behavior of `dbm_store` in case a matching record already exists in the database. It can have one of the following two values:

`DBM_REPLACE`

Replace existing record with the new one.

`DBM_INSERT`

The existing record is left unchanged, and the function returns `'1'`.

If no matching record exists in the database, new record will be inserted no matter what the value of the `mode` is.

`int dbm_delete (DBM *dbf, datum key)` [ndbm]

Deletes the record with the matching key from the database. If the function succeeds, `'0'` is returned. Otherwise, if no matching record is found or if an error occurs, `'-1'` is returned.

`datum dbm_firstkey (DBM *dbf)` [ndbm]

Initializes iteration over the keys from the database and returns the first key. Note, that the word `'first'` does not imply any specific ordering of the keys.

If there are no records in the database, the `dptr` member of the returned datum is `'NULL'`. Otherwise, the `dptr` member of the returned datum points to the memory managed by the compatibility library. The application should never free it.

`datum dbm_nextkey (DBM *dbf)` [ndbm]

Continues the iteration started by `dbm_firstkey`. Returns the next key in the database. If the iteration covered all keys in the database, the `dptr` member of the returned datum is `'NULL'`. Otherwise, the `dptr` member of the returned datum points to the memory managed by the compatibility library. The application should never free it.

The usual way of iterating over all the records in the database is:

```
for (key = dbm_firstkey (dbf);
     key.ptr;
     key = dbm_nextkey (dbf))
{
    /* do something with the key */
}
```

The loop above should not try to delete any records from the database, otherwise the iteration is not guaranteed to cover all the keys. See [Chapter 10 \[Sequential\]](#), page 10, for a detailed discussion of this.

**int dbm\_error** (*DBM \*dbf*) [ndbm]  
 Returns the error condition of the database: '0' if no errors occurred so far while manipulating the database, and a non-zero value otherwise.

**void dbm\_clearerr** (*DBM \*dbf*) [ndbm]  
 Clears the error condition of the database.

**int dbm\_dirfno** (*DBM \*dbf*) [ndbm]  
 Returns the file descriptor of the 'dir' file of the database. It is guaranteed to be different from the descriptor returned by the `dbm_pagfno` function (see below).  
 The application can lock this descriptor to serialize accesses to the database.

**int dbm\_pagfno** (*DBM \*dbf*) [ndbm]  
 Returns the file descriptor of the 'pag' file of the database. See also `dbm_dirfno`.

**int dbm\_rdonly** (*DBM \*dbf*) [ndbm]  
 Returns '1' if the database *dbf* is open in a read-only mode and '0' otherwise.

## 19.2 DBM interface functions.

The functions below are provided for compatibility with the old UNIX 'DBM' interface. Only one database at a time can be manipulated using them.

**int dbmopen** (*char \*file*) [dbm]  
 Opens a database. The *file* argument is the full name of the database file to be opened. The function opens two files: *file.pag* and *file.dir*. If any of them does not exist, the function fails. It never attempts to create the files.  
 The database is opened in the read-write mode, if its disk permissions permit.  
 The application must ensure that the functions described below in this section are called only after a successful call to `dbmopen`.

**int dbmclose** (*void*) [dbm]  
 Closes the database opened by an earlier call to `dbmopen`.

**datum fetch** (*datum key*) [dbm]  
 Reads a record from the database with the matching key. The *key* argument supplies the key that is being looked for.  
 If no matching record is found, the `dptr` member of the returned datum is 'NULL'. Otherwise, the `dptr` member of the returned datum points to the memory managed by the compatibility library. The application should never free it.

**int store** (*datum key, datum content*) [dbm]  
 Stores the key/value pair in the database. If a record with the matching key already exists, its content will be replaced with the new one.  
 Returns '0' on success and '-1' on error.

**int delete** (*datum key*) [dbm]  
 Deletes a record with the matching key.  
 If the function succeeds, '0' is returned. Otherwise, if no matching record is found or if an error occurs, '-1' is returned.

`datum firstkey (void)` [dbm]

Initializes iteration over the keys from the database and returns the first key. Note, that the word ‘**first**’ does not imply any specific ordering of the keys.

If there are no records in the database, the `dptr` member of the returned datum is ‘NULL’. Otherwise, the `dptr` member of the returned datum points to the memory managed by the compatibility library. The application should never free it.

`datum nextkey (datum key)` [dbm]

Continues the iteration started by a call to `firstkey`. Returns the next key in the database. If the iteration covered all keys in the database, the `dptr` member of the returned datum is ‘NULL’. Otherwise, the `dptr` member of the returned datum points to the memory managed by the compatibility library. The application should never free it.

## 20 Examine and modify a GDBM database.

The `gdbmtool` utility allows you to view and modify an existing GDBM database or to create a new one.

When invoked without arguments, it tries to open a database file called `junk.gdbm`, located in the current working directory. You can change this default by supplying the name of the database to use as an argument to the program, e.g.:

```
$ gdbmtool file.db
```

The database will be opened in read-write mode, unless the `-r` (`--read-only`) option is specified, in which case it will be opened only for reading.

If the database does not exist, `gdbmtool` will create it. There is a special option `-n` (`--newdb`, which instructs the utility to create a new database. If it is used and if the database already exists, it will be deleted, so use it sparingly.

### 20.1 gdbmtool invocation

The following table summarizes all `gdbmtool` command line options:

<code>-b size</code>	
<code>--block-size=size</code>	Set block size.
<code>-c size</code>	
<code>--cache-size=size</code>	Set cache size.
<code>-f file</code>	
<code>--file file</code>	Read commands from <i>file</i> , instead of the standard input.
<code>-h</code>	
<code>--help</code>	Print a concise help summary.
<code>-N</code>	
<code>--norc</code>	Don't read startup files (see <a href="#">Section 20.2.4 [startup files]</a> , page 40).
<code>-n</code>	
<code>--newdb</code>	Create the database.
<code>-l</code>	
<code>--no-lock</code>	Disable file locking.
<code>-m</code>	
<code>--no-mmap</code>	Disable mmap.
<code>-q</code>	
<code>--quiet</code>	Don't print the usual welcome banner at startup. This is the same as setting the variable 'quiet' in the startup file. See <a href="#">[quiet]</a> , page 35.

```

-r
--read-only
    Open the database in read-only mode.

-s
--synchronize
    Synchronize to the disk after each write.

-V
--version
    Print program version and licensing information and exit.

--usage
    Print a terse invocation syntax summary along with a list of available command
    line options.

```

## 20.2 gdbmtool interactive mode

After successful startup, `gdbmtool` starts a loop, in which it reads commands from the standard input, executes them and prints the results on the standard output. If the standard input is attached to a console, `gdbmtool` runs in interactive mode, which is indicated by its *prompt*:

```
gdbmtool> _
```

The utility finishes when it reads the ‘quit’ command (see below) or detects end-of-file on its standard input, whichever occurs first.

A `gdbmtool` command consists of a *command verb*, optionally followed by *arguments*, separated by any amount of white space. A command verb can be entered either in full or in an abbreviated form, as long as that abbreviation does not match any other verb. For example, ‘co’ can be used instead of ‘count’ and ‘ca’ instead of ‘cache’.

Any sequence of non-whitespace characters appearing after the command verb forms an argument. If the argument contains whitespace or unprintable characters it must be enclosed in double quotes. Within double quotes the usual *escape sequences* are understood, as shown in the table below:

Sequence	Replaced with
<code>\a</code>	Audible bell character (ASCII 7)
<code>\b</code>	Backspace character (ASCII 8)
<code>\f</code>	Form-feed character (ASCII 12)
<code>\n</code>	Newline character (ASCII 10)
<code>\r</code>	Carriage return character (ASCII 13)
<code>\t</code>	Horizontal tabulation character (ASCII 9)
<code>\v</code>	Vertical tabulation character (ASCII 11)
<code>\\</code>	Single slash
<code>\"</code>	Double quote

Table 20.1: Backslash escapes

In addition, a backslash immediately followed by the end-of-line character effectively removes that character, allowing to split long arguments over several input lines.

Command parameters may be optional or mandatory. If the number of actual arguments is less than the number of mandatory parameters, `gdbmtool` will prompt you to supply

missing arguments. For example, the ‘**store**’ command takes two mandatory parameters, so if you invoked it with no arguments, you would be prompted twice to supply the necessary data, as shown in example below:

```
gdbmtool> store
key? three
data? 3
```

However, such prompting is possible only in interactive mode. In non-interactive mode (e.g. when running a script), all arguments must be supplied with each command, otherwise **gdbmtool** will report an error and exit immediately.

### 20.2.1 Shell Variables

A number of **gdbmtool** parameters is kept in its internal variables.

**bool confirm** [gdbmtool variable]

Whether to ask for confirmation before certain destructive operations, such as truncating the existing database.

Default is ‘**true**’.

**string ps1** [gdbmtool variable]

Primary prompt string. Its value can contain *conversion specifiers*, consisting of the ‘%’ character followed by another character. These specifiers are expanded in the resulting prompt as follows:

Sequence	Expansion
%f	name of the current database file
%p	program invocation name
%P	package name (‘GDBM’)
%v	program version
%_	single space character
%%	%

The default value is ‘%p>%\_’, i.e. the program name, followed by a “greater than” sign, followed by a single space.

**string ps2** [gdbmtool variable]

Secondary prompt. See ‘**ps1**’ for a description of its value. This prompt is displayed before reading the second and subsequent lines of a multi-line command.

The default value is ‘%\_>%\_’.

**string delim1** [gdbmtool variable]

A string used to delimit fields of a structured datum on output (see [Section 20.2.3 \[definitions\]](#), page 39).

Default is ‘,’ (a comma). This variable cannot be unset.

**string delim2** [gdbmtool variable]

A string used to delimit array items when printing a structured datum (see [Section 20.2.3 \[definitions\]](#), page 39).

Default is ‘,’ (a comma). This variable cannot be unset.

**string pager** [gdbmtool variable]

The name and command line of the pager program to pipe output to. This program is used in interactive mode when the estimated number of output lines is greater than the number of lines on your screen.

The default value is inherited from the environment variable `PAGER`. Unsetting this variable disables paging.

**bool quiet** [gdbmtool variable]

Whether to display a welcome banner at startup. This variable should be set in a startup script file (see [Section 20.2.4 \[startup files\]](#), page 40). See [\[-q option\]](#), page 32.

The following variables control how the database is opened:

**numeric blocksize** [gdbmtool variable]

Sets the block size. See [Chapter 4 \[Open\]](#), page 4. Unset by default.

**numeric cachesize** [gdbmtool variable]

Sets the cache size. See [Chapter 15 \[Options\]](#), page 19. By default this variable is not set.

**string open** [gdbmtool variable]

Open mode. The following values are allowed:

**newdb** Truncate the database if it exists or create a new one. Open it in read-write mode.

Technically, this sets the ‘GDBM\_NEWDB’ flag in call to ‘gdbm\_open’. See [Chapter 4 \[Open\]](#), page 4.

**wrcreat**

**rw** Open the database in read-write mode. Create it if it does not exist. This is the default.

Technically speaking, it sets the ‘GDBM\_WRCREAT’ flag in call to `gdbm_open`. See [Chapter 4 \[Open\]](#), page 4.

**reader**

**readonly** Open the database in read-only mode. Signal an error if it does not exist.

This sets the ‘GDBM\_READER’ flag (see [Chapter 4 \[Open\]](#), page 4).

Attempting to set any other value or to unset this variable produces an error.

**number filemode** [gdbmtool variable]

File mode (in octal) for creating new database files and database dumps.

**bool lock** [gdbmtool variable]

Lock the database. This is the default.

Setting this variable to false or unsetting it results in passing ‘GDBM\_NOLOCK’ flag to `gdbm_open` (see [Chapter 4 \[Open\]](#), page 4).

**bool mmap** [gdbmtool variable]

Use memory mapping. This is the default.

Setting this variable to false or unsetting it results in passing ‘GDBM\_NOMMAP’ flag to `gdbm_open` (see [Chapter 4 \[Open\]](#), page 4).



**bool sync** [gdbmtool variable]  
 Flush all database writes on disk immediately. Default is false. See [Chapter 4 \[Open\]](#), page 4.

The following commands are used to list or modify the variables:

**set [assignments]** [command verb]  
 When used without arguments, lists all variables and their values. Unset variables are shown after a comment sign ('#'). For string and numeric variables, values are shown after an equals sign. For boolean variables, only the variable name is displayed if the variable is 'true'. If it is 'false', its name is prefixed with 'no'.

For example:

```
ps1="%p>%_"
ps2="%_>%_"
delim1=","
delim2=";"
confirm
# cachesize is unset
# blocksize is unset
open="wrcreat"
lock
mmap
nosync
pager="less"
# quiet is unset
```

If used with arguments, the **set** command alters the specified variables. In this case, arguments are variable assignments in the form '**name=value**'. For boolean variables, the *value* is interpreted as follows: if it is numeric, '0' stands for 'false', any non-zero value stands for 'true'. Otherwise, the values 'on', 'true', and 'yes' denote 'true', and 'off', 'false', 'no' stand for 'false'. Alternatively, only the name of a boolean variable can be supplied to set it to 'true', and its name prefixed with 'no' can be used to set it to false. For example, the following command sets the 'delim2' variable to ';' and the 'confirm' variable to 'false':

```
set delim2=";" noconfirm
```

**unset variables** [command verb]  
 Unsets the listed variables. The effect of unsetting depends on the variable. Unless explicitly described in the discussion of the variables above, unsetting a boolean variable is equivalent to setting it to 'false'. Unsetting a string variable is equivalent to assigning it an empty string.

## 20.2.2 Gdbmtool Commands

**avail** [command verb]  
 Print the *avail list*.

**bucket num** [command verb]  
 Print the bucket number *num* and set it as the current one.

<b>cache</b>	[command verb]
Print the bucket cache.	
<b>close</b>	[command verb]
Close the currently open database.	
<b>count</b>	[command verb]
Print the number of entries in the database.	
<b>current</b>	[command verb]
Print the current bucket.	
<b>delete key</b>	[command verb]
Delete record with the given <i>key</i>	
<b>dir</b>	[command verb]
Print hash directory.	
<b>export file-name</b> [ <i>truncate</i> ] [ <i>binary ascii</i> ]	[command verb]
Export the database to the flat file <i>file-name</i> . See <a href="#">Chapter 13 [Flat files]</a> , <a href="#">page 14</a> , for a description of the flat file format and its purposes. This command will not overwrite an existing file, unless the ‘ <i>truncate</i> ’ parameter is also given. Another optional argument determines the type of the dump (see <a href="#">Chapter 13 [Flat files]</a> , <a href="#">page 14</a> ). By default, ASCII dump is created.	
The global variable <b>filemode</b> specifies the permissions to use for the created output file.	
See also <a href="#">Chapter 23 [gdbmexport]</a> , <a href="#">page 44</a> .	
<b>fetch key</b>	[command verb]
Fetch and display the record with the given <i>key</i> .	
<b>first</b>	[command verb]
Fetch and display the first record in the database. Subsequent records can be fetched using the <b>next</b> command (see below). See <a href="#">Chapter 10 [Sequential]</a> , <a href="#">page 10</a> , for more information on sequential access.	
<b>hash key</b>	[command verb]
Compute and display the hash value for the given <i>key</i> .	
<b>header</b>	[command verb]
Print file header.	
<b>help</b>	[command verb]
<b>?</b>	[command verb]
Print a concise command summary, showing each command verb with its parameters and a short description of what it does. Optional arguments are enclosed in square brackets.	

**import** *file-name* [*replace*] [*nometa*] [command verb]

Import data from a flat dump file *file-name* (see [Chapter 13 \[Flat files\]](#), page 14). If the word ‘**replace**’ is given as an argument, any records with the same keys as the already existing ones will replace them. The word ‘**nometa**’ turns off restoring meta-information from the dump file.

**list** [command verb]

List the contents of the database.

**next** [*key*] [command verb]

Sequential access: fetch and display the next record. If the *key* is given, the record following the one with this key will be fetched.

See also **first**, above.

See [Chapter 10 \[Sequential\]](#), page 10, for more information on sequential access.

**open** *filename* [command verb]

Open the database file *filename*. If successful, any previously open database is closed. Otherwise, if the operation fails, the currently opened database remains unchanged.

This command takes additional information from the following variables:

‘**open**’        The database access mode. See [\[The open variable\]](#), page 35, for a list of its values.

‘**lock**’        Whether or not to lock the database. Default is ‘**on**’.

‘**mmap**’        Use the memory mapping. Default is ‘**on**’.

‘**sync**’        Synchronize after each write. Default is ‘**off**’.

‘**filemode**’        Specifies the permissions to use in case a new file is created.

See [\[open parameters\]](#), page 35, for a detailed description of these variables.

**quit** [command verb]

Close the database and quit the utility.

**reorganize** [command verb]

Reorganize the database (see [Chapter 11 \[Reorganization\]](#), page 12).

**source** *filename* [command verb]

Read `gdbmtool` commands from the file *filename*.

**status** [command verb]

Print current program status. The following example shows the information displayed:

```
Database file: junk.gdbm
Database is open
define key string
define content string
```

The two ‘**define**’ strings show the defined formats for key and content data. See [Section 20.2.3 \[definitions\]](#), page 39, for a detailed discussion of their meaning.

**store key data** [command verb]

Store the *data* with *key* in the database. If *key* already exists, its data will be replaced.

**version** [command verb]

Print the version of *gdbm*.

### 20.2.3 Data Definitions

GDBM databases are able to keep data of any type, both in the key and in the content part of a record. Quite often these data are structured, i.e. they consist of several fields of various types. *Gdbmtool* provides a mechanism for handling such kind of records.

The **define** command defines a record structure. The general syntax is:

```
define what definition
```

where *what* is ‘**key**’ to defining the structure of key data and ‘**content**’ to define the structure of the content records.

The *definition* can be of two distinct formats. In the simplest case it is a single data type. For example,

```
define content int
```

defines content records consisting of a single integer field. Supported data types are:

char	Single byte (signed).
short	Signed short integer.
ushort	Unsigned short integer.
int	Signed integer.
unsigned	
uint	Unsigned integer.
long	Signed long integer.
ulong	Unsigned long integer.
llong	Signed long long integer.
ullong	Unsigned long long integer.
float	A floating point number.
double	Double-precision floating point number.
string	Array of bytes.
stringz	Null-terminated string, trailing null being part of the string.

All numeric data types (integer as well as floating point) have the same respective widths as in C language on the host where the database file resides.

The ‘**string**’ and ‘**stringz**’ are special. Both define a string of bytes, similar to ‘**char x[]**’ in C. The former defines an array of bytes, the latter - a null-terminated string. This makes a difference, in particular, when the string is the only part of datum. Consider the following two definitions:

1. **define key string**

## 2. `define` key stringz

Now, suppose we want to store the string "ab" in the key. Using the definition (1), the `dptr` member of GDBM datum will contain two bytes: 'a', and 'b'. Consequently, the `dsize` member will have the value 2. Using the definition (2), the `dptr` member will contain three bytes: 'a', 'b', and ASCII 0. The `dsize` member will have the value 3.

The definition (1) is the default for both key and content.

The second form of the `define` statement is similar to the C `struct` statement and allows for defining structural data. In this form, the *definition* part is a comma-separated list of data types and variables enclosed in curly braces. In contrast to the rest of `gdbm` commands, this command is inherently multiline and is terminated with the closing curly brace. For example:

```
define content {
    int status,
    pad 8,
    char id[3],
    string name
}
```

This defines a structure consisting of three members: an integer `status`, an array of 8 bytes `id`, and a null-terminated string `name`. Notice the `pad` statement: it allows to introduce padding between structure members. Another useful statement is `offset`: it specifies that the member following it begins at the given offset in the structure. Assuming the size of `int` is 8 bytes, the above definition can also be written as

```
define content {
    int status,
    offset 16,
    char id[3],
    string name
}
```

*NOTE:* The 'string' type can reasonably be used only if it is the last or the only member of the data structure. That's because it provides no information about the number of elements in the array, so it is interpreted to contain all bytes up to the end of the datum.

When displaying the structured data, `gdbmtool` precedes each value with the corresponding field name and delimits parts of the structure with the string defined in the 'delim1' variable (see [Section 20.2.1 \[variables\]](#), page 34). Array elements are delimited using the string from 'delim2'. For example:

```
gdbmtool> fetch foo
status=2,id={ a, u, x },name="quux"
```

To supply a structured datum as an argument to a `gdbmtool` command, use the same notation, but without field names, e.g.:

```
gdbmtool> hash { 2, {a,u,x}, "quux" }
hash value = 13089969.
```

### 20.2.4 Startup Files

Upon startup `gdbmtool` looks for a file named '`.gdbmtoolrc`' first in the current working directory and, if not found, in the home directory of the user who started the command.

The port compiled with DJGPP will look either for a file named `‘.gdbmtoolrc’` or `‘_gdbmtoolrc’`, in that order, depending on if the underlying file system used offers long filename support (aka LFN support) or not. If not, like on plain DOS, then only for `‘_gdbmtoolrc’` will be looked.

If found, this file is read and interpreted as a list of `gdbmtool` commands. This allows you to customize the program behavior.

Following is an example startup file which disables the welcome banner, sets command line prompt to contain the name of the database file in parentheses and defines the structure of the database content records:

```
set quiet
set ps1="( %f) "
define key stringz
define content {
    int time,
    pad 4,
    int status
}
```

## 21 The `gdbm_dump` utility

The `gdbm_dump` utility creates a flat file dump of a GDBM database (see [Chapter 13 \[Flat files\]](#), [page 14](#)). It takes one mandatory argument: the name of the source database file. The second argument, if given, specifies the name of the output file. If not given, `gdbm_dump` will produce the dump on the standard output.

For example, the following invocation creates a dump of the database `file.db` in the file `file.dump`:

```
$ gdbm_dump file.db file.dump
```

By default the utility creates dumps in ASCII format (see [Chapter 13 \[Flat files\]](#), [page 14](#)). Another format can be requested using the `--format` (`-H`) option.

The `gdbm_dump` utility understands the following command line options:

`-H fmt`

`--format=fmt`

Select output format. Valid values for *fmt* are: ‘binary’ or ‘0’ to select binary dump format, and ‘ascii’ or ‘1’ to select ASCII format.

`-h`

`--help` Print a concise help summary.

`-V`

`--version`

Print program version and licensing information and exit.

`--usage`

Print a terse invocation syntax summary along with a list of available command line options.

## 22 The `gdbm_load` utility

The `gdbm_load` utility restores a GDBM database from a flat file. The utility requires at least one argument: the name of the input flat file. If it is '-', the standard input will be read. The format of the input file is detected automatically.

By default the utility attempts to restore the database under its original name, as stored in the input file. It will fail to do so if the input is in binary format. In that case, the name of the database must be given as the second argument.

In general, if two arguments are given the second one is treated as the name of the database to create, overriding the file name specified in the flat file.

The utility understands the following command line arguments:

```
-b num
--block-size=num
    Sets block size. See Chapter 4 \[Open\], page 4.

-c num
--cache-size=num
    Sets cache size. See Chapter 15 \[Options\], page 19.

-M
--mmap    Use memory mapping.

-m mode
--mode=mode
    Sets the file mode. The argument is the desired file mode in octal.

-n
--no-meta
    Do not restore file meta-data (ownership and mode) from the flat file.

-r
--replace
    Replace existing keys.

-u user[:group]
--user=user[:group]
    Set file owner. The user can be either a valid user name or UID. Similarly,
    the group is either a valid group name or GID. If group is not given, the main
    group of user is used.
    User and group parts can be separated by a dot, instead of the colon.

-h
--help    Print a concise help summary.

-V
--version
    Print program version and licensing information and exit.

--usage   Print a terse invocation syntax summary along with a list of available command
line options.
```



## 23 Export a database into a portable format.

The `gdbmexport` utility converts the database of an older GDBM version into a binary flat format.

The utility takes two mandatory arguments: the name of the database file to convert and the output file name, e.g.:

```
$ gdbmexport junk.gdbm junk.flat
```

In addition the following two options are understood:

- `-h`        Display short usage summary and exit.
- `-v`        Display program version and licensing information, and exit.

## 24 Exit codes

All GDBM utilities return uniform exit codes. These are summarized in the table below:

<b>Code</b>	<b>Meaning</b>
0	Successful termination.
1	A fatal error occurred.
2	Program was unable to restore file ownership or mode.
3	Command line usage error.

## 25 Problems and bugs.

If you have problems with GNU `dbm` or think you've found a bug, please report it. Before reporting a bug, make sure you've actually found a real bug. Carefully reread the documentation and see if it really says you can do what you're trying to do. If it's not clear whether you should be able to do something or not, report that too; it's a bug in the documentation!

Before reporting a bug or trying to fix it yourself, try to isolate it to the smallest possible input file that reproduces the problem. Then send us the input file and the exact results `gdbm` gave you. Also say what you expected to occur; this will help us decide whether the problem was really in the documentation.

Once you've got a precise problem, send e-mail to [bug-gdbm@gnu.org](mailto:bug-gdbm@gnu.org).

Please include the version number of GNU `dbm` you are using. You can get this information by printing the variable `gdbm_version` (see [Chapter 17 \[Variables\]](#), page 23).

Non-bug suggestions are always welcome as well. If you have questions about things that are unclear in the documentation or are just obscure features, please report them too.

You may contact the authors and maintainers by e-mail:

[phil@cs.wvu.edu](mailto:phil@cs.wvu.edu), [downsj@downsj.com](mailto:downsj@downsj.com), [gray@gnu.org.ua](mailto:gray@gnu.org.ua)

## 26 Additional resources

For the latest updates and pointers to additional resources, visit <http://www.gnu.org/software/gdbm>.

In particular, a copy of `gdbm` documentation in various formats is available online at <http://www.gnu.org/software/gdbm/manual.html>.

Latest versions of `gdbm` can be downloaded from anonymous FTP: <ftp://ftp.gnu.org/gnu/gdbm>, or via HTTP from <http://ftp.gnu.org/gnu/gdbm>, or from any GNU mirror worldwide. See <http://www.gnu.org/order/ftp.html>, for a list of mirrors.

To track `gdbm` development, visit <http://puszcza.gnu.org.ua/projects/gdbm>.

# Appendix A GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright © 2000, 2001, 2002, 2007, 2008, 2011 Free Software  
Foundation, Inc.  
<http://fsf.org/>

Everyone is permitted to copy and distribute verbatim copies  
of this license document, but changing it is not allowed.

## 0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document *free* in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

## 1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “publisher” means any person or entity that distributes copies of the Document to the public.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

## 2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

## 3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

## 4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at



your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

## 5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements."

## 6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

## 7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

## 8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

## 9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

## 10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

## 11. RELICENSING

“Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

## ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

```
Copyright (C)  year  your name.
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.3
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover
Texts. A copy of the license is included in the section entitled ‘‘GNU
Free Documentation License’’.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

```
with the Invariant Sections being list their titles, with
the Front-Cover Texts being list, and with the Back-Cover Texts
being list.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

# Index

-		
--newdb, gdbmtool option	32	
--read-only, gdbmtool option	32	
-n, gdbmtool option	32	
-r, gdbmtool option	32	
.		
.gdbmtoolrc	40	
?		
? .....	37	
-		
_GDBM_MAX_ERRNO	23	
_GDBM_MIN_ERRNO	23	
A		
avail	36	
B		
blocksize	35	
bucket	36	
C		
cache	37	
cachesize	35	
close	37	
close-on-exec	4	
closing database	5	
command line options, gdbmtool	32	
compatibility layer	28	
confirm	34	
count	37	
creating a database, gdbmtool	32	
current	37	
D		
database options	19	
database reorganization	12	
database synchronization	13	
database, closing	5	
database, opening or creating	4	
DBM functions	30	
dbm.h	28	
dbm_clearerr	30	
dbm_close	28	
dbm_delete	29	
dbm_dirfno	30	
dbm_error	30	
dbm_fetch	28	
dbm_firstkey	29	
DBM_INSERT	29	
dbm_nextkey	29	
dbm_open	28	
dbm_pagfno	30	
dbm_rdonly	30	
DBM_REPLACE	29	
dbm_store	29	
dbmclose	30	
dbmopen	30	
default database, gdbmtool	32	
delete	30, 37	
deleting records	9	
deletion in iteration loops	10	
delim1	34	
delim2	34	
dir	37	
'dir' file	28	
E		
error codes	25	
error strings	18	
exit code	45	
export	14	
export	37	
F		
fetch	30, 37	
fetching records	8	
filemode	35	
first	37	
firstkey	31	
Flat file format	14	
G		
GDBM_BAD_FILE_OFFSET	26	
GDBM_BAD_MAGIC_NUMBER	25	
GDBM_BAD_OPEN_FLAGS	26	
GDBM_BLOCK_SIZE_ERROR	25	
GDBM_BYTE_SWAPPED	26	
GDBM_CACHESIZE	19	
GDBM_CANNOT_REPLACE	26	
GDBM_CANT_BE_READER	25	
GDBM_CANT_BE_WRITER	25	
GDBM_CENTFREE	20	
GDBM_CLOEXEC	4	
gdbm_close	5	
GDBM_COALESCEBLKS	20	
gdbm_count	6	

gdbm\_delete..... 9  
 gdbm\_delete and sequential access..... 10  
 gdbm\_dump..... 14, 42  
 gdbm\_dump\_to\_file..... 16  
 GDBM\_EMPTY\_DATABASE..... 25  
 GDBM\_ERR\_FILE\_MODE..... 15, 27  
 GDBM\_ERR\_FILE\_OWNER..... 15, 27  
 gdbm\_errlist[]..... 23  
 gdbm\_errno..... 23  
 gdbm\_exists..... 8  
 gdbm\_export..... 16  
 gdbm\_export\_to\_file..... 16  
 GDBM\_FASTMODE..... 19  
 gdbm\_fdesc..... 22  
 gdbm\_fetch..... 8  
 GDBM\_FILE\_EOF..... 27  
 GDBM\_FILE\_OPEN\_ERROR..... 25  
 GDBM\_FILE\_READ\_ERROR..... 25  
 GDBM\_FILE\_SEEK\_ERROR..... 25  
 GDBM\_FILE\_STAT\_ERROR..... 26  
 GDBM\_FILE\_WRITE\_ERROR..... 25  
 gdbm\_firstkey..... 10  
 GDBM\_GETCACHESIZE..... 19  
 GDBM\_GETCOALESCEBLKS..... 20  
 GDBM\_GETDBNAME..... 20  
 GDBM\_GETFLAGS..... 19  
 GDBM\_GETMAXMAPSIZE..... 20  
 GDBM\_GETMMAP..... 20  
 GDBM\_GETSYNCMODE..... 20  
 GDBM\_ILLEGAL\_DATA..... 26  
 gdbm\_import..... 16  
 gdbm\_import\_from\_file..... 16  
 GDBM\_INSERT..... 7  
 GDBM\_ITEM\_NOT\_FOUND..... 26  
 gdbm\_load..... 15, 43  
 gdbm\_load\_from\_file..... 16  
 GDBM\_MALLOC\_ERROR..... 25  
 GDBM\_NEWDB..... 4  
 gdbm\_nextkey..... 10  
 GDBM\_NO\_DBNAME..... 27  
 GDBM\_NO\_ERROR..... 25  
 GDBM\_NOLOCK..... 4, 22  
 GDBM\_NOMMAP..... 4  
 gdbm\_open..... 4  
 GDBM\_OPT\_ALREADY\_SET..... 26  
 GDBM\_OPT\_ILLEGAL..... 26  
 GDBM\_READER..... 4  
 GDBM\_READER\_CANT\_DELETE..... 25  
 GDBM\_READER\_CANT\_REORGANIZE..... 26  
 GDBM\_READER\_CANT\_STORE..... 26  
 gdbm\_reorganize..... 12  
 GDBM\_REORGANIZE\_FAILED..... 26  
 GDBM\_REPLACE..... 7  
 GDBM\_SETCACHESIZE..... 19  
 GDBM\_SETCENTFREE..... 20  
 GDBM\_SETCOALESCEBLKS..... 20  
 GDBM\_SETMAXMAPSIZE..... 20  
 GDBM\_SETMMAP..... 20

gdbm\_setopt..... 19  
 GDBM\_SETSYNCMODE..... 19  
 gdbm\_store..... 7  
 gdbm\_strerror..... 18  
 gdbm\_sync..... 13  
 GDBM\_SYNC..... 4, 13  
 GDBM\_SYNCMODE..... 19  
 GDBM\_UNKNOWN\_UPDATE..... 26  
 gdbm\_version..... 23  
 gdbm\_version\_cmp..... 23  
 GDBM\_VERSION\_MAJOR..... 23  
 GDBM\_VERSION\_MINOR..... 23  
 gdbm\_version\_number[3]..... 23  
 GDBM\_VERSION\_PATCH..... 23  
 GDBM\_WRCREAT..... 4  
 GDBM\_WRITER..... 4  
 gdbmexport..... 44  
 gdbmtool..... 32

## H

hash..... 37  
 header..... 37  
 help..... 37

## I

import..... 14  
 import..... 38  
 init file, gdbmtool..... 40  
 interactive mode, gdbmtool..... 33  
 iterating over records..... 10  
 iteration and gdbm\_delete..... 10  
 iteration loop..... 10  
 iteration loop, using 'NDBM'..... 29

## J

junk.gdbm..... 32

## L

libgdbm\_compat..... 28  
 list..... 38  
 lock..... 35  
 locking..... 22  
 looking up records..... 8

## M

mmap..... 35

## N

NDBM functions..... 28  
 ndbm.h..... 28  
 next..... 38

**nextkey** ..... 31  
**number of records** ..... 6

## O

**open** ..... 35, 38  
**opening the database** ..... 4  
**options, database** ..... 19

## P

**'pag' file** ..... 28  
**pager** ..... 35  
**ps1** ..... 34  
**ps2** ..... 34

## Q

**quiet** ..... 35  
**quit** ..... 38

## R

**read-only mode, gdbmtool** ..... 32  
**record, deleting** ..... 9  
**record, fetching** ..... 8  
**records, iterating over** ..... 10

**records, storing** ..... 7  
**records, testing existence** ..... 8  
**reorganization, database** ..... 12  
**reorganize** ..... 38

## S

**sequential access** ..... 10  
**sequential access, using 'NDBM'** ..... 29  
**set** ..... 36  
**source** ..... 38  
**startup file, gdbmtool** ..... 40  
**status** ..... 38  
**store** ..... 30, 39  
**storing records** ..... 7  
**sync** ..... 36  
**synchronization, database** ..... 13

## U

**unset** ..... 36

## V

**variables, gdbmtool** ..... 34  
**version** ..... 39  
**version number** ..... 23