

**NAME**

lmdb – LMDB Lightning Memory-Mapped Database binding for gawk

**SYNOPSIS**

```
@load "lmdb"
```

```
handle = mdb_env_create()
```

**DESCRIPTION**

The *lmdb* extension implements the mdb C API. Please refer to the `<lmdb.h>` header file for documentation, or view it on the web at <http://www.lmdb.tech/doc>. This document explains only the peculiarities of the *gawk* implementation. Aside from the `mdb_strerror()` function, every function in this library should set the `MDB_ERRNO` variable before returning. So you can always check whether (`MDB_ERRNO == MDB_SUCCESS`) to see if the last call succeeded. The library also sets the standard *gawk* `ERRNO` variable if an error occurs.

Almost all of the `mdb_*` functions in `<lmdb.h>` are implemented, as of version 0.9.18.

To achieve optimal performance, please use cursors whenever possible. Under the hood, the mdb library always uses cursors.

**Global Variables****MDB\_SUCCESS**

0, as defined in `<lmdb.h>`.

**MDB\_ERRNO**

Return status from the last mdb function call.

**MDB\_KEY, MDB\_DATA**

Subscripts for use with `mdb_cursor_get()`.

**MDB** An array defining all the `MDB_*` constants in `<lmdb.h>` with the "MDB\_" prefix removed. For example, `MDB["FIRST"]` contains the value of `MDB_FIRST`. The various flags arguments can be constructed by using the *gawk* `or()` function to combine flag bits defined in this array.

**Functions****string mdb\_strerror(<int errno>)**

This function takes an integer argument and returns the result from calling the `mdb_strerror()` library function. The *gawk* library defines an additional extension-specific error code that is also supported by this function. If the argument is not an integer, it returns an empty string.

**string mdb\_env\_create()**

Call `mdb_env_create()` and return a string handle to use for referencing this environment. On error, an empty string is returned, and `MDB_ERRNO` contains the return code from `mdb_env_create()`.

**int mdb\_env\_open(<env handle>, <path>, <u\_int flags>, <u\_int mode>)**

Returns the status, which can also be found in `MDB_ERRNO`. You can use the *gawk* `or()` function to construct the flags. For example, the third argument might be `or(MDB["NOSUBDIR"], MDB["NOSYNC"], MDB["NOLOCK"])`.

**int mdb\_env\_close(<env handle>)**

Call `mdb_env_close()` and return status, which can also be found in `MDB_ERRNO`. Although the C function has no return value, this function can return an error if the handle argument is invalid. If the result is `MDB_SUCCESS`, the handle is released and may no longer be used.

**int mdb\_env\_sync(<env handle>, <int force>)**

Returns the status, which can also be found in `MDB_ERRNO`.

**int mdb\_env\_copy(<env handle>, <path string>)**

Returns the status, which can also be found in `MDB_ERRNO`.

- int mdb\_env\_copy2(<env handle>, <path string>, <u\_int flags>)**  
Returns the status, which can also be found in **MDB\_ERRNO**.
- int mdb\_env\_get\_flags(<env handle>)**  
On error, returns 0, else the flags value. Please compare **MDB\_ERRNO** to **MDB\_SUCCESS** to see whether the call succeeded.
- int mdb\_env\_get\_maxkeysize(<env handle>)**  
On error, returns 0, else the maxkeysize value. Please compare **MDB\_ERRNO** to **MDB\_SUCCESS** to see whether the call succeeded.
- int mdb\_env\_get\_maxreaders(<env handle>)**  
On error, returns 0, else the maxreaders value. Please compare **MDB\_ERRNO** to **MDB\_SUCCESS** to see whether the call succeeded.
- string mdb\_env\_get\_path(<env handle>)**  
On error, returns "", else the path value. Please compare **MDB\_ERRNO** to **MDB\_SUCCESS** to see whether the call succeeded.
- int mdb\_env\_set\_flags(<env handle>, <u\_int flags>, <int onoff>)**  
Returns the status, which can also be found in **MDB\_ERRNO**.
- int mdb\_env\_set\_mapsize(<env handle>, <u\_int mapsize>)**  
Returns the status, which can also be found in **MDB\_ERRNO**.
- int mdb\_env\_set\_maxdbs(<env handle>, <u\_int dbs>)**  
Returns the status, which can also be found in **MDB\_ERRNO**.
- int mdb\_env\_set\_maxreaders(<env handle>, <u\_int readers>)**  
Returns the status, which can also be found in **MDB\_ERRNO**.
- string mdb\_txn\_begin(<env handle>, <parent txn handle or empty string>, <u\_int flags>)**  
Call *mdb\_txn\_begin()* and return a string handle to use for referencing this transaction. On error, an empty string is returned, and **MDB\_ERRNO** will contain an error code.
- int mdb\_txn\_id(<txn handle>)**  
On error, returns 0, else the transaction id. Please compare **MDB\_ERRNO** to **MDB\_SUCCESS** to see whether the call succeeded.
- int mdb\_txn\_commit(<txn handle>)**  
Returns the status, which can also be found in **MDB\_ERRNO**. If the result is **MDB\_SUCCESS**, the handle is released and may no longer be used.
- int mdb\_txn\_abort(<txn handle>)**  
Returns the status, which can also be found in **MDB\_ERRNO**. If the result is **MDB\_SUCCESS**, the handle is released and may no longer be used.
- int mdb\_txn\_reset(<txn handle>)**  
Returns the status, which can also be found in **MDB\_ERRNO**.
- int mdb\_txn\_renew(<txn handle>)**  
Returns the status, which can also be found in **MDB\_ERRNO**.
- string mdb\_dbi\_open(<txn handle>, <database name or empty string>, <u\_int flags>)**  
Call *mdb\_dbi\_open()* and return a string handle to use for referencing this transaction. On error, an empty string is returned, and **MDB\_ERRNO** will contain an error code.
- int mdb\_dbi\_close(<env handle>, <dbi handle>)**  
Returns the status, which can also be found in **MDB\_ERRNO**. If the result is **MDB\_SUCCESS**, the handle is released and may no longer be used.
- int mdb\_dbi\_flags(<txn handle>, <dbi handle>)**  
On error, returns 0, else the flags value. Please compare **MDB\_ERRNO** to **MDB\_SUCCESS** to see whether the call succeeded.

**int mdb\_drop(<txn handle>, <dbi handle>, <del: 0 or 1>)**

Returns the status, which can also be found in **MDB\_ERRNO**. If the call succeeds and <del> is 1, the dbi handle is released and may no longer be used.

**string mdb\_get(<txn handle>, <dbi handle>, <key>)**

On error, returns "", else the data string returned by the function. You must check **MDB\_ERRNO** to see whether the call succeeded, since an empty string could be a valid data value.

Please note that starting with version 2 of the API available in *gawk* 4.2, the result will be flagged as user input data eligible for the `strnum` attribute. So if the value appears to be numeric, *gawk* will treat it as a number in comparisons. This feature was not available prior to version 2 of the *gawk* API.

**int mdb\_put(<txn handle>, <dbi handle>, <key>, <data>, <u\_int flags>)**

Returns the status, which can also be found in **MDB\_ERRNO**.

**int mdb\_del(<txn handle>, <dbi handle>, <key>[, <data>])**

Returns the status, which can also be found in **MDB\_ERRNO**. If the fourth argument is not present, a NULL pointer is passed as the fourth argument to the underlying function. Please be careful: passing an empty "" string as the fourth argument is not the same as omitting the fourth argument!

**string mdb\_cursor\_open(<txn handle>, <dbi handle>)**

Call *mdb\_cursor\_open()* and return a string handle to use for referencing this transaction. On error, an empty string is returned, and **MDB\_ERRNO** will contain an error code.

**int mdb\_cursor\_close(<cursor handle>)**

Returns the status, which can also be found in **MDB\_ERRNO**.

**int mdb\_cursor\_renew(<cursor handle>)**

Returns the status, which can also be found in **MDB\_ERRNO**.

**int mdb\_cursor\_get(<cursor handle>, <key/data array>, <u\_int cursor op>)**

Returns the status, which can also be found in **MDB\_ERRNO**. The second argument is an array. If the **MDB\_KEY** element is populated, its value is passed as the second argument to the underlying function. In its absence, an empty "" string is passed. Similarly, the contents of the **MDB\_DATA** element are used to construct the third argument to the underlying function. On exit, if the function call succeeds, these same array elements will be populated with the data returned.

Please note that starting with version 2 of the API available in *gawk* 4.2, the returned array elements will be flagged as user input data eligible for the `strnum` attribute. So if the key or data appears to be numeric, *gawk* will treat it as a number in comparisons. This feature was not available prior to version 2 of the *gawk* API.

**int mdb\_cursor\_put(<cursor handle>, <key>, <data>, <u\_int flags>)**

Returns the status, which can also be found in **MDB\_ERRNO**.

**int mdb\_cursor\_del(<cursor handle>, <u\_int flags>)**

Returns the status, which can also be found in **MDB\_ERRNO**.

**int mdb\_cursor\_count(<cursor handle>)**

On error, returns 0, else the function result. Please compare **MDB\_ERRNO** to **MDB\_SUCCESS** to see whether the call succeeded.

**int mdb\_reader\_check(<env handle>)**

On error, returns 0, else the dead value from the function. Please compare **MDB\_ERRNO** to **MDB\_SUCCESS** to see whether the call succeeded.

**int mdb\_cmp(<txn handle>, <dbi handle>, <string a>, <string b>)**

On error, returns 0, else the result of the underlying function. **MDB\_ERRNO** to **MDB\_SUCCESS** to see whether the call succeeded.

**int mdb\_dcmp(<txn handle>, <dbi handle>, <string a>, <string b>)**

On error, returns 0, else the result of the underlying function. **MDB\_ERRNO** to **MDB\_SUCCESS** to see whether the call succeeded.

**int mdb\_env\_stat(<env handle>, <array>)**

Returns the status, which can also be found in **MDB\_ERRNO**. If the call succeeds, the array is cleared and populated with the elements of the **MDB\_stat** structure stripped of the "ms\_" prefix.

**int mdb\_stat(<txn handle>, <dbi handle>, <array>)**

Returns the status, which can also be found in **MDB\_ERRNO**. If the call succeeds, the array is cleared and populated with the elements of the **MDB\_stat** structure stripped of the "ms\_" prefix.

**int mdb\_env\_info(<env handle>, <array>)**

Returns the status, which can also be found in **MDB\_ERRNO**. If the call succeeds, the array is cleared and populated with the elements of the **MDB\_envinfo** structure stripped of the "me\_" prefix.

**string mdb\_txn\_env(<txn handle>)**

On error, an empty "" string is returned, else the env handle associated with this transaction.

**string mdb\_cursor\_txn(<cursor handle>)**

On error, an empty "" string is returned, else the transaction handle associated with this cursor.

**string mdb\_cursor\_dbi(<cursor handle>)**

On error, an empty "" string is returned, else the dbi handle associated with this cursor.

**string mdb\_version([<version array>])**

Call *mdb\_version()* and return the version string. If the optional array argument is supplied, it will be cleared, and the **major**, **minor**, and **patch** subscripts will be populated with the associated version numbers.

The following functions are not implemented: *mdb\_env\_copyfd*, *mdb\_env\_copyfd2*, *mdb\_env\_get\_fd*, *mdb\_env\_get\_userctx*, *mdb\_env\_set\_assert*, *mdb\_env\_set\_userctx*, *mdb\_reader\_list*, *mdb\_set\_compare*, *mdb\_set\_dupsort*, *mdb\_set\_relctx*, *mdb\_set\_relfunc*.

Of these, *mdb\_set\_compare* and *mdb\_set\_dupsort* seem most important. To implement these, we need a way for the extension to call into a *gawk* function. I do not think that is currently possible.

## EXAMPLE

Please refer to **dict.awk** and **dict\_cursor.awk** located in the **test** directory.

## SEE ALSO

**/usr/include/lmdb.h**, <http://www.lmdb.tech/doc>

## AUTHOR

Andrew Schorr

## COPYING PERMISSIONS

Copyright © 2016, Free Software Foundation, Inc.

Permission is granted to make and distribute verbatim copies of this manual page provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this manual page under the conditions for verbatim copying, provided that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this manual page into another language, under the above conditions for modified versions, except that this permission notice may be stated in a translation approved by the Foundation.