Full credit is given to the above companies including the OS that this PDF file was generated!

Red Hat Enterprise Linux Release 9.2 Manual Pages on 'mlocate.db.5' command

$ man mlocate.db.5

mlocate.db(5) File Formats Manual mlocate.db(5)

NAME

mlocate.db - a mlocate database

DESCRIPTION

A mlocate database starts with a file header: 8 bytes for a magic num?
ber ("\0mlocate" like a C literal), 4 bytes for the configuration block
size in big endian, 1 byte for file format version (0), 1 byte for the
?require visibility? flag (0 or 1), 2 bytes padding, and a NUL-termi?
nated path name of the root of the database.

The header is followed by a configuration block, included to ensure
databases are not reused if some configuration changes could affect
their contents. The size of the configuration block in bytes is stored
in the file header. The configuration block is a sequence of variable
assignments, ordered by variable name. Each variable assignment con?
sists of a NUL-terminated variable name and an ordered list of NUL-ter?
minated values. The value list is terminated by one more NUL charac?
ter. The ordering used is defined by the strcmp () function.

Currently defined variables are:

prune_bind_mounts

A single entry, the value of PRUNE_BIND_MOUNTS; one of the strings 0 or 1.

prunefs

The value of PRUNEFS, each entry is converted to uppercase.

prunepaths

The value of PRUNEPATHS.

The rest of the file until EOF describes directories and their con? tents. Each directory starts with a header: 8 bytes for directory time (seconds) in big endian, 4 bytes for directory time (nanoseconds) in big endian (0 if unknown, less than 1,000,000,000), 4 bytes padding, and a NUL-terminated path name of the the directory. Directory con? tents, a sequence of file entries sorted by name, follow.

Directory time is the maximum of st_ctime and st_mtime of the direc? tory. updatedb(8) uses the original data if the directory time in the database and in the file system match exactly. Directory time equal to 0 always causes rescanning of the directory: this is necessary to han? dle directories which were being updated while building the database.

Each file entry starts with a single byte, marking its type:

0 A non-directory file. Followed by a NUL-terminated file (not path) name.

1 A subdirectory. Followed by a NUL-terminated file (not path) name.

2 Marks the end of the current directory.

locate(1) only reports file entries, directory names are not reported because they are reported as an entry in their parent directory. The only exception is the root directory of the database, which is stored in the file header.

AUTHOR

Miloslav Trmac <mitr@redhat.com>

SEE ALSO

locate(1), updatedb.conf(5), updatedb(8)