NAME

dmplfs − dump logical file system to tape

SYNOPSIS

dmplfs lfs_name tape_unit#

DESCRIPTION

Dmplfs copies a logical file system (LFS) to tape. Unlike *dd*, which does a device-to-device copy, *dmplfs* writes two tape files, the first containing the LFS overhead area (header, file definition entries, freelist and bitmap), and the second the contents of all allocated logical files in ascending order. *Dmplfs* can be used to save the contents of the LFS for later restoral by *rstlfs*(1), or, because of the manner in which the files are dumped and restored, to compress an LFS whose free space has become highly fragmented (freelist full of areas too small to be used).

Lfs_name is the filename of the LFS in /dev and *tape_unit#* is the number of the tape drive on which the dump tape is mounted. Both parameters are required, and the program assumes that the tape is 2400 ft. long and will be written at 1600 bpi. For convenience, the user may specify the tape unit as 0-3; the program will modify the unit number as necessary to get the correct density. If the command is entered with no parameters, the program will print the expected syntax.

Dmplfs assumes that the overhead file will fit on one tape reel and that the data file may require more than one reel; the program will prompt the user when a new reel is to be mounted. An 80-character label file is written at the beginning of each reel (including the first) which contains the *lfs_name*, reel number, date and time. The blocking factor for both the overhead and data files is 5120 bytes (10 sectors) per tape block. In the overhead file, all tape blocks are full size (5120) except possibly the last block in the file, which may be shorter. In the data file, every logical file is written beginning on a tape block boundary, and if the file is less than 10 sectors long the tape block contains only the allocated file size in units of LF blocks. Similarly, the last tape block of a logical file contains only the remainder of the file in units of LF blocks.

FILES

/dev/*lfs_name*	LFS to be written to tape
/dev/mt*tape_unit#*	tape unit to be used
/etc/lmtab	list of mounted logical file systems

SEE ALSO

lfcheck(1), mklfs(1), rstlfs(1)

DIAGNOSTICS

Dmplfs prints self-explanatory error messages on exit whenever a problem is detected.

WARNINGS

Dmplfs uses the start and size information in the file definition entries to read the logical files from disk which can result in the "unfolding" of overlapped files (files containing duplicated blocks) as well as attempts to read overhead or bad blocks which have been erroneously allocated to files. These side effects can be prevented by making sure that the LFS checks (using *lfcheck*(1)) before dumping to tape.

Do not attempt to dump a mounted logical filesystem; the LFS should be unmounted and flushed to disk before *dmplfs*
is invoked.

The LFS should be re-made using *mklfs*(1) before restoring with *rstlfs*(1). As additional insurance, it is wise to make a dd tape of the LFS block device before doing the *mklfs* so the LFS can be restored to its prior state if necessary (i.e., if *rstlfs* has trouble reading the *dmplfs* tape).

Dmplfs assumes that the 1600 bpi tape units have file names **/dev/mt8** − **/dev/mt11** (rewind) and **/dev/mt12** − **/dev/mt15** (no rewind).

BUGS

In order to prevent the tape running off the end of the reel, there is an artificial limit of 6400 tape blocks per reel for the data file. This number was chosen to allow room for the drive to skip over bad spots on the tape when writing.