

**NAME**

**lockf** - execute a command while holding a file lock

**SYNOPSIS**

**lockf** [-kns~~w~~] [-t *seconds*] *file command* [*arguments*]

**DESCRIPTION**

The **lockf** utility acquires an exclusive lock on a *file*, creating it if necessary, *and removing the file on exit unless explicitly told not to*. While holding the lock, it executes a *command* with optional *arguments*. After the *command* completes, **lockf** releases the lock, and removes the *file* unless the **-k** option is specified. BSD-style locking is used, as described in flock(2); the mere existence of the *file* is not considered to constitute a lock.

If the **lockf** utility is being used to facilitate concurrency between a number of processes, it is recommended that the **-k** option be used. This will guarantee lock ordering, as well as implement a performance enhanced algorithm which minimizes CPU load associated with concurrent unlink, drop and re-acquire activity. It should be noted that if the **-k** option is not used, then no guarantees around lock ordering can be made.

The following options are supported:

- k** Causes the lock file to be kept (not removed) after the command completes.
- s** Causes **lockf** to operate silently. Failure to acquire the lock is indicated only in the exit status.
- n** Causes **lockf** to fail if the specified lock *file* does not exist. If **-n** is not specified, **lockf** will create *file* if necessary.
- t *seconds*** Specifies a timeout for waiting for the lock. By default, **lockf** waits indefinitely to acquire the lock. If a timeout is specified with this option, **lockf** will wait at most the given number of *seconds* before giving up. A timeout of 0 may be given, in which case **lockf** will fail unless it can acquire the lock immediately. When a lock times out, *command* is *not* executed.
- w** Causes **lockf** to open *file* for writing rather than reading. This is necessary on filesystems (including NFSv4) where a file which has been opened read-only cannot be exclusively locked.

In no event will **lockf** break a lock that is held by another process.

## EXIT STATUS

If **lockf** successfully acquires the lock, it returns the exit status produced by *command*. Otherwise, it returns one of the exit codes defined in `sysexits(3)`, as follows:

**EX\_TEMPFAIL** The specified lock file was already locked by another process.

**EX\_CANTCREAT** The **lockf** utility was unable to create the lock file, e.g., because of insufficient access privileges.

**EX\_UNAVAILABLE**  
The **-n** option is specified and the specified lock file does not exist.

**EX\_USAGE** There was an error on the **lockf** command line.

**EX\_OSERR** A system call (e.g., `fork(2)`) failed unexpectedly.

**EX\_SOFTWARE** The *command* did not exit normally, but may have been signaled or stopped.

## EXAMPLES

The first job takes a lock and sleeps for 5 seconds in the background. The second job tries to get the lock and timeouts after 1 second (PID numbers will differ):

```
$ lockf mylock sleep 5 & lockf -t 1 mylock echo "Success"
[1] 94410
lockf: mylock: already locked
```

The first job takes a lock and sleeps for 1 second in the background. The second job waits up to 5 seconds to take the lock and echoes the message on success (PID numbers will differ):

```
$ lockf mylock sleep 1 & lockf -t 5 mylock echo "Success"
[1] 19995
Success
[1]+ Done          lockf mylock sleep 1
```

## SEE ALSO

`flock(2)`, `lockf(3)`, `sysexits(3)`

## HISTORY

A **lockf** utility first appeared in FreeBSD 2.2.

**AUTHORS**

John Polstra <[jdp@polstra.com](mailto:jdp@polstra.com)>