

fdopen(3)

NAME  
fdopen – associate a stream with a file descriptor

SYNOPSIS  
#include <stdio.h>  
FILE \*fdopen(int *fdes*, const char \**mode*);

DESCRIPTION  
The **fdopen()** function associates a stream with a file descriptor *fdes*, whose value must be less than 255. The *mode* argument is a character string having one of the following values:  
**r** or **rb** open a file for reading  
**w** or **wb** open a file for writing  
**a** or **ab** open a file for writing at end of file  
**r+** or **rb+** or **r+b** open a file for update (reading and writing)  
**w+** or **wb+** or **w+b** open a file for update (reading and writing)  
**a+** or **ab+** or **a+b** open a file for update (reading and writing) at end of file  
The meaning of these flags is exactly as specified in **fopen(3S)**, except that modes beginning with **w** do not cause truncation of the file.  
The mode of the stream must be allowed by the file access mode of the open file. The file position indicator associated with the new stream is set to the position indicated by the file offset associated with the file descriptor.  
**fdopen()** will preserve the offset maximum previously set for the open file description corresponding to *fdes*.  
The error and end-of-file indicators for the stream are cleared. The **fdopen()** function may cause the **st\_atime** field of the underlying file to be marked for update.  
RETURN VALUES  
Upon successful completion, **fdopen()** returns a pointer to a stream. Otherwise, a null pointer is returned and **errno** is set to indicate the error.  
**fdopen()** may fail and not set **errno** if there are no free **stdio** streams.  
ERRORS  
The **fdopen()** function may fail if:  
**EBADF** The *fdes* argument is not a valid file descriptor.  
**EINVAL** The *mode* argument is not a valid mode.  
**EMFILE** **FOPEN\_MAX** streams are currently open in the calling process.  
**EMFILE** **STREAM\_MAX** streams are currently open in the calling process.  
**ENOMEM** Insufficient space to allocate a buffer.  
USAGE  
**STREAM\_MAX** is the number of streams that one process can have open at one time. If defined, it has the same value as **FOPEN\_MAX**.  
File descriptors are obtained from calls like **open(2)**, **dup(2)**, **creat(2)** or **pipe(2)**, which open files but do not return streams. Streams are necessary input for almost all of the Section 3S library routines.  
SEE ALSO  
**creat(2)**, **dup(2)**, **open(2)**, **pipe(2)**, **fclose(3S)**, **fopen(3S)**, **attributes(5)**

accept(2)

NAME  
accept – accept a connection on a socket

SYNOPSIS  
#include <sys/types.h>  
#include <sys/socket.h>  
int accept(int *s*, struct sockaddr \**addr*, int \**addrlen*);

DESCRIPTION  
The argument *s* is a socket that has been created with **socket(3N)** and bound to an address with **bind(3N)**, and that is listening for connections after a call to **listen(3N)**. The **accept()** function extracts the first connection on the queue of pending connections, creates a new socket with the properties of *s*, and allocates a new file descriptor, *ns*, for the socket. If no pending connections are present on the queue and the socket is not marked as non-blocking, **accept()** blocks the caller until a connection is present. If the socket is marked as non-blocking and no pending connections are present on the queue, **accept()** returns an error as described below. The **accept()** function uses the **netconfig(4)** file to determine the STREAMS device file name associated with *s*. This is the device on which the connect indication will be accepted. The accepted socket, *ns*, is used to read and write data to and from the socket that connected to *ns*; it is not used to accept more connections. The original socket (*s*) remains open for accepting further connections.  
The argument *addr* is a result parameter that is filled in with the address of the connecting entity as it is known to the communications layer. The exact format of the *addr* parameter is determined by the domain in which the communication occurs.  
The argument *addrlen* is a value-result parameter. Initially, it contains the amount of space pointed to by *addr*; on return it contains the length in bytes of the address returned.  
The **accept()** function is used with connection-based socket types, currently with **SOCK\_STREAM**.  
It is possible to **select(3C)** or **poll(2)** a socket for the purpose of an **accept()** by selecting or polling it for a read. However, this will only indicate when a connect indication is pending; it is still necessary to call **accept()**.  
RETURN VALUE  
On success, these system calls return a nonnegative integer that is a file descriptor for the accepted socket. On error, **-1** is returned, and *errno* is set appropriately.  
ERRORS  
**accept()** will fail if:  
**EBADF** The descriptor is invalid.  
**EINTR** The accept attempt was interrupted by the delivery of a signal.  
**EMFILE** The per-process descriptor table is full.  
**ENODEV** The protocol family and type corresponding to *s* could not be found in the **netconfig** file.  
**ENOMEM** There was insufficient user memory available to complete the operation.  
**EPROTO** A protocol error has occurred; for example, the STREAMS protocol stack has not been initialized or the connection has already been released.  
**EWOULDBLOCK** The socket is marked as non-blocking and no connections are present to be accepted.  
SEE ALSO  
**poll(2)**, **bind(3N)**, **connect(3N)**, **listen(3N)**, **select(3C)**, **socket(3N)**, **netconfig(4)**, **attributes(5)**, **socket(5)**

**NAME**  
 fnmatch – match filename or pathname

**SYNOPSIS**  

```
#include <fnmatch.h>

int fnmatch(const char *pattern, const char *string, int flags);
```

**DESCRIPTION**  
 The `fnmatch()` function checks whether the *string* argument matches the *pattern* argument, which is a shell wildcard pattern.

The *flags* argument modifies the behavior; it is the bitwise OR of zero or more of the following flags:

**FNM\_NOESCAPE**  
 If this flag is set, treat backslash as an ordinary character, instead of an escape character.

**FNM\_PATHNAME**  
 If this flag is set, match a slash in *string* only with a slash in *pattern* and not by an asterisk (\*) or a question mark (?) metacharacter, not by a bracket expression ([]) containing a slash.

**FNM\_PERIOD**  
 If this flag is set, a leading period in *string* has to be matched exactly by a period in *pattern*. A period is considered to be leading if it is the first character in *string*, or if both **FNM\_PATHNAME** and **FNM\_FILE\_NAME** is set and the period immediately follows a slash.

**FNM\_FILE\_NAME**  
 This is a GNU synonym for **FNM\_PATHNAME**.

**FNM\_LEADING\_DIR**  
 If this flag (a GNU extension) is set, the pattern is considered to be matched if it matches an initial segment of *string* which is followed by a slash. This flag is mainly for the internal use of `glibc` and is only implemented in certain cases.

**FNM\_CASEFOLD**  
 If this flag (a GNU extension) is set, the pattern is matched case-insensitively.

**RETURN VALUE**  
 Zero if *string* matches *pattern*, **FNM\_NOMATCH** if there is no match or another nonzero value if there is an error.

**CONFORMING TO**  
 POSIX.2. The **FNM\_FILE\_NAME**, **FNM\_LEADING\_DIR**, and **FNM\_CASEFOLD** flags are GNU extensions.

**NAME**  
 fflush – flush a stream

**SYNOPSIS**  

```
#include <stdio.h>

int fflush(FILE *stream);
```

**DESCRIPTION**  
 For output streams, `fflush()` forces a write of all user-space buffered data for the given output or update *stream* via the stream's underlying write function.

For input streams associated with seekable files (e.g., disk files, but not pipes or terminals), `fflush()` discards any buffered data that has been fetched from the underlying file, but has not been consumed by the application.

The open status of the stream is unaffected.

If the *stream* argument is `NULL`, `fflush()` flushes *all* open output streams.

For a nonlocking counterpart, see `unlocked_stdio(3)`.

**RETURN VALUE**  
 Upon successful completion 0 is returned. Otherwise, **EOF** is returned and *errno* is set to indicate the error.

**ERRORS**  
**EBADF**  
*stream* is not an open stream, or is not open for writing.

The function `fflush()` may also fail and set *errno* for any of the errors specified for `write(2)`.

**SEE ALSO**  
`fsync(2)`, `sync(2)`, `write(2)`, `fclose(3)`, `fileno(3)`, `fopen(3)`, `setbuf(3)`, `unlocked_stdio(3)`



**NAME** pthread\_create – create a new thread / pthread\_exit – terminate the calling thread

**SYNOPSIS**  
#include <pthread.h>

```
int pthread_create(pthread_t * thread, pthread_attr_t * attr, void * (*start_routine)(void *), void *
arg);
void pthread_exit(void *retval);
```

**DESCRIPTION**

**pthread\_create** creates a new thread of control that executes concurrently with the calling thread. The new thread applies the function *start\_routine* passing it *arg* as first argument. The new thread terminates either explicitly, by calling **pthread\_exit(3)**, or implicitly, by returning from the *start\_routine* function. The latter case is equivalent to calling **pthread\_exit(3)** with the result returned by *start\_routine* as exit code.

The *attr* argument specifies thread attributes to be applied to the new thread. See **pthread\_attr\_init(3)** for a complete list of thread attributes. The *attr* argument can also be **NULL**, in which case default attributes are used; the created thread is joinable (not detached) and has default (non real-time) scheduling policy.

**pthread\_exit** terminates the execution of the calling thread. All cleanup handlers that have been set for the calling thread with **pthread\_cleanup\_push(3)** are executed in reverse order (the most recently pushed handler is executed first). Finalization functions for thread-specific data are then called for all keys that have non-**NULL** values associated with them in the calling thread (see **pthread\_key\_create(3)**). Finally, execution of the calling thread is stopped.

The *retval* argument is the return value of the thread. It can be consulted from another thread using **pthread\_join(3)**.

**RETURN VALUE**

On success, the identifier of the newly created thread is stored in the location pointed by the *thread* argument, and a 0 is returned. On error, a non-zero error code is returned.

The **pthread\_exit** function never returns.

**ERRORS**

**EAGAIN**

not enough system resources to create a process for the new thread.

**EAGAIN**

more than **PTHREAD\_THREADS\_MAX** threads are already active.

**AUTHOR**

Xavier Leroy <Xavier.Leroy@inria.fr>

**SEE ALSO**

**pthread\_join(3)**, **pthread\_detach(3)**, **pthread\_attr\_init(3)**.

**NAME**

sigaction – POSIX signal handling functions.

**SYNOPSIS**

#include <signal.h>

```
int sigaction(int signal, const struct sigaction *act, struct sigaction *oldact);
```

**DESCRIPTION**

The **sigaction** system call is used to change the action taken by a process on receipt of a specific signal.

*signal* specifies the signal and can be any valid signal except **SIGKILL** and **SIGSTOP**.

If *act* is non-null, the new action for signal *signal* is installed from *act*. If *oldact* is non-null, the previous action is saved in *oldact*.

The **sigaction** structure is defined as something like

```
struct sigaction {
    void (*sa_handler)(int signal_number);
    sigset_t sa_mask;
    int sa_flags;
}
```

*sa\_handler* specifies the action to be associated with *signal* and may be **SIG\_DFL** for the default action, **SIG\_IGN** to ignore this signal, or a pointer to a signal handling function.

*sa\_mask* gives a mask of signals which should be blocked during execution of the signal handler. In addition, the signal which triggered the handler will be blocked, unless the **SA\_NODEFER** or **SA\_NOMASK** flags are used.

*sa\_flags* specifies a set of flags which modify the behaviour of the signal handling process. It is formed by the bitwise OR of zero or more of the following:

**SA\_NOCLDSTOP**

If *signal* is **SIGCHLD**, do not receive notification when child processes stop (i.e., when child processes receive one of **SIGSTOP**, **SIGTSTP**, **SIGTTIN** or **SIGTTOU**).

**SA\_RESTART**

Provide behaviour compatible with BSD signal semantics by making certain system calls restartable across signals. Without **SA\_RESTART** the system calls return an error and set *errno* to **EINTR** when interrupted by a signal.

**RETURN VALUES**

**sigaction(0)** returns 0 on success; on error, **-1** is returned, and *errno* is set to indicate the error.

**ERRORS**

**EINVAL**

An invalid signal was specified. This will also be generated if an attempt is made to change the action for **SIGKILL** or **SIGSTOP**, which cannot be caught.

**SEE ALSO**

**kill(1)**, **kill(2)**, **killpg(2)**, **pause(2)**, **sigsetops(3)**.

**NAME**

stat, fstat, lstat — get file status

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/stat.h>
#include <unistd.h>

int stat(const char *path, struct stat *buf);
int fstat(int fd, struct stat *buf);
int lstat(const char *path, struct stat *buf);
```

Feature Test Macro Requirements for glibc (see [feature\\_test\\_macros\(7\)](#)):

```
lstat(_BSD_SOURCE || _XOPEN_SOURCE >= 500
```

**DESCRIPTION**

These functions return information about a file. No permissions are required on the file itself, but — in the case of `stat()` and `lstat()` — execute (search) permission is required on all of the directories in `path` that lead to the file.

`stat()` stats the file pointed to by `path` and fills in `buf`.

`lstat()` is identical to `stat()`, except that if `path` is a symbolic link, then the link itself is stat-ed, not the file that it refers to.

`fstat()` is identical to `stat()`, except that the file to be stat-ed is specified by the file descriptor `fd`.

All of these system calls return a `stat` structure, which contains the following fields:

```
struct stat {
    dev_t   st_dev;      /* ID of device containing file */
    ino_t   st_ino;     /* inode number */
    mode_t  st_mode;    /* protection */
    nlink_t st_nlink;   /* number of hard links */
    uid_t   st_uid;     /* user ID of owner */
    gid_t   st_gid;     /* group ID of owner */
    dev_t   st_rdev;    /* device ID (if special file) */
    off_t   st_size;    /* total size, in bytes */
    blksize_t st_blksize; /* blocksize for file system I/O */
    blkcnt_t st_blocks; /* number of blocks allocated */
    time_t  st_atime;   /* time of last access */
    time_t  st_mtime;   /* time of last modification */
    time_t  st_ctime;   /* time of last status change */
};
```

The `st_dev` field describes the device on which this file resides.

The `st_rdev` field describes the device that this file (inode) represents.

The `st_size` field gives the size of the file (if it is a regular file or a symbolic link) in bytes. The size of a symlink is the length of the pathname it contains, without a trailing null byte.

The `st_blocks` field indicates the number of blocks allocated to the file, 512-byte units. (This may be smaller than `st_size/512` when the file has holes.)

The `st_blksize` field gives the "preferred" blocksize for efficient file system I/O. (Writing to a file in smaller chunks may cause an inefficient read-modify-rewrite.)

Not all of the Linux file systems implement all of the time fields. Some file system types allow mounting in such a way that file accesses do not cause an update of the `st_atime` field. (See "noatime" in [mount\(8\)](#).)

The field `st_atime` is changed by file accesses, for example, by `execve(2)`, `mknod(2)`, `pipe(2)`, `utime(2)` and `read(2)` (of more than zero bytes). Other routines, like `mmap(2)`, may or may not update `st_atime`.

The field `st_mtime` is changed by file modifications, for example, by `mknod(2)`, `truncate(2)`, `utime(2)` and `write(2)` (of more than zero bytes). Moreover, `st_mtime` of a directory is changed by the creation or deletion of files in that directory. The `st_mtime` field is *not* changed for changes in owner, group, hard link count, or mode.

The field `st_ctime` is changed by writing or by setting inode information (i.e., owner, group, link count, mode, etc.).

The following POSIX macros are defined to check the file type using the `st_mode` field:

```
S_ISREG(m)    is it a regular file?
S_ISDIR(m)    directory?
S_ISCHR(m)    character device?
S_ISBLK(m)    block device?
S_ISFIFO(m)   FIFO (named pipe)?
S_ISLNK(m)    symbolic link? (Not in POSIX.1-1996.)
S_ISSOCK(m)   socket? (Not in POSIX.1-1996.)
```

**RETURN VALUE**

On success, zero is returned. On error, `-1` is returned, and `errno` is set appropriately.

**ERRORS**

```
EACCES      Search permission is denied for one of the directories in the path prefix of path. (See also path\_resolution\(7\).)
EBADF      fd is bad.
EFAULT     Bad address.
ELOOP      Too many symbolic links encountered while traversing the path.
ENAMETOOLONG File name too long.
ENOENT     A component of the path path does not exist, or the path is an empty string.
ENOMEM     Out of memory (i.e., kernel memory).
ENOTDIR    A component of the path is not a directory.
```

**SEE ALSO**

[access\(2\)](#), [chmod\(2\)](#), [chown\(2\)](#), [fstatat\(2\)](#), [readlink\(2\)](#), [utime\(2\)](#), [capabilities\(7\)](#), [symlink\(7\)](#)

