

fopen/fdopen/fileno(3)

fopen/fdopen/fileno(3)

getc/fgets/putc/fputs(3)

getc/fgets/putc/fputs(3)

NAME

fopen, fdopen, fileno – stream open functions

NAME

fgetc, fgets, getc, getchar, fputc, fputs, putc, putchar – input and output of characters and strings

SYNOPSIS

#include <stdio.h>

SYNOPSIS

#include <stdio.h>

FILE *fopen(const char *path, const char *mode);
FILE *fdopen(int fdidx, const char *mode);
int fileno(FILE *stream);

int fgetc(FILE *stream);
char *fgets(char *, int size, FILE *stream);
int getc(FILE *stream);
int getchar(void);
int fputc(int c, FILE *stream);
int fputs(const char *, FILE *stream);
int puts(int c, FILE *stream);
int putchar(int c);

DESCRIPTION

The **fopen** function opens the file whose name is the string pointed to by *path* and associates a stream with it.
The argument *mode* points to a string beginning with one of the following sequences (Additional characters may follow these sequences.):

- r** Open text file for reading. The stream is positioned at the beginning of the file.
- r+** Open for reading and writing. The stream is positioned at the beginning of the file.
- w** Truncate file to zero length or create text file for writing. The stream is positioned at the beginning of the file.
- w+** Open for reading and writing. The file is created if it does not exist, otherwise it is truncated. The stream is positioned at the beginning of the file.
- a** Open for appending (writing at end of file). The file is created if it does not exist. The stream is positioned at the end of the file.
- a+** Open for reading and appending (writing at end of file). The file is created if it does not exist. The stream is positioned at the end of the file.

The **fdopen** function associates a stream with the existing file descriptor, *fdidx*. The *mode* of the stream (one of the values "r", "r+", "w", "w+", "a", "a+") must be compatible with the mode of the file descriptor. The file position indicator of the new stream is set to that belonging to *fdidx*, and the error and end-of-file indicators are cleared. Modes "w" or "w+" do not cause truncation of the file. The file descriptor is not dup'ed, and will be closed when the stream created by **fdopen** is closed. The result of applying **fdopen** to a shared memory object is undefined.

The function **fileno**(3) examines the argument *stream* and returns its integer descriptor.

RETURN VALUE

Upon successful completion **fopen**, **fdopen** and **freopen** return a **FILE** pointer. Otherwise, **NULL** is returned and the global variable *errno* is set to indicate the error.

ERRORS

EINVAL

The *mode* provided to **fopen**, **fdopen**, or **freopen** was invalid.

The **fopen**, **fdopen** and **freopen** functions may also fail and set *errno* for any of the errors specified for the routine **malloc**(3).

The **fopen** function may also fail and set *errno* for any of the errors specified for the routine **open**(2).

The **fdopen** function may also fail and set *errno* for any of the errors specified for the routine **fcntl**(2).

SEE ALSO

open(2), **fclose**(3), **fileno**(3)

DESCRIPTION

fgetc(3) reads the next character from *stream* and returns it as an *unsigned char* cast to an *int*, or **EOF** on end of file or error.

getc(3) is equivalent to **fgetc**(3) except that it may be implemented as a macro which evaluates *stream* more than once.

getchar(3) is equivalent to **getc**(3).

fgets(3) reads in at most one less than *size* characters from *stream* and stores them into the buffer pointed to by *s*. Reading stops after an **EOF** or a newline. If a newline is read, it is stored into the buffer. A '\0' is stored after the last character in the buffer.

fputc(3) writes the character *c*, cast to an *unsigned char*, to *stream*.

fputs(3) writes the string *s* to *stream*, without its terminating null byte ('\0').

putc(3) is equivalent to **fputc**(3) except that it may be implemented as a macro which evaluates *stream* more than once.

putchar(3); is equivalent to **putc**(3, *stdout*).

Calls to the functions described here can be mixed with each other and with calls to other output functions from the *stdio* library for the same output stream.

RETURN VALUE

fgetc(3), **getc**(3) and **getchar**(3) return the character read as an *unsigned char* cast to an *int* or **EOF** on end of file or error.

fgets(3) returns *s* on success, and **NULL** on error or when end of file occurs while no characters have been read. **fputc**(3), **putc**(3) and **putchar**(3) return the character written as an *unsigned char* cast to an *int* or **EOF** on error.

fputs(3) returns a nonnegative number on success, or **EOF** on error.

SEE ALSO

read(2), **write**(2), **fcntl**(2), **ferror**(3), **feof**(3), **fdopen**(3), **fdopen**(3), **fgetc**(3), **fgetc**(3), **fseek**(3), **getwchar**(3), **scanf**(3), **ungetc**(3), **write**(2), **ferror**(3), **fopen**(3), **fopen**(3), **fputc**(3), **fputc**(3), **fseek**(3), **fwrite**(3), **gets**(3), **putwchar**(3), **scanf**(3), **unlock**(3), **unlocked_stdio**(3)

<p>opendir/readdir(3)</p> <p>NAME opendir – open a directory / readdir – read a directory to the directory stream.</p> <p>SYNOPSIS #include <sys/types.h> #include <dirent.h></p> <p>DIR *opendir(const char *name); struct dirent *readdir(DIR *dir);</p> <p>DESCRIPTION opendir The opendir() function opens a directory stream corresponding to the directory <i>name</i>, and returns a pointer to the directory stream. The stream is positioned at the first entry in the directory.</p> <p>RETURN VALUE The opendir() function returns a pointer to the directory stream or NULL if an error occurred.</p> <p>DESCRIPTION readdir The readdir() function returns a pointer to a dirent structure representing the next directory entry in the directory stream pointed to by <i>dir</i>. It returns NULL on reaching the end-of-file or if an error occurred. It is safe to use readdir() inside threads if the pointers passed as <i>dir</i> are created by distinct calls to opendir(). The data returned by readdir() is overwritten by subsequent calls to readdir() for the same directory stream.</p> <p>The <i>dirent</i> structure is defined as follows:</p> <pre> struct dirent { long d_ino; /* inode number */ char d_name[256]; /* filename */ }; </pre> <p>RETURN VALUE The readdir() function returns a pointer to a dirent structure, or NULL if an error occurs or end-of-file is reached.</p> <p>ERRORS EACCES Permission denied. ENOENT Directory does not exist, or <i>name</i> is an empty string. ENOTDIR <i>name</i> is not a directory.</p>	<p>opendir/readdir(3)</p>	<p>stat(2)</p> <p>NAME stat, lstat – get file status</p> <p>SYNOPSIS #include <sys/types.h> #include <sys/stat.h> #include <unistd.h></p> <p>int stat(const char *path, struct stat *buf); int lstat(int fd, struct stat *buf); int lstat(const char *path, struct stat *buf);</p> <p>Feature Test Macro Requirements for glibc (see feature_test_macros(7)):</p> <p>lstat(): _BSD_SOURCE _XOPEN_SOURCE >= 500</p> <p>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 <i>path</i> that lead to the file.</p> <p>stat() stats the file pointed to by <i>path</i> and fills in <i>buf</i>.</p> <p>lstat() is identical to stat(), except that if <i>path</i> is a symbolic link, then the link itself is stat-ed, not the file that it refers to.</p> <p>fstat() is identical to stat(), except that the file to be stat-ed is specified by the file descriptor <i>fd</i>.</p> <p>All of these system calls return a <i>stat</i> structure, which contains the following fields:</p> <pre> 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 */ }; </pre> <p>The <i>st_dev</i> field describes the device on which this file resides.</p> <p>The <i>st_rdev</i> field describes the device that this file (inode) represents.</p> <p>The <i>st_size</i> 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.</p> <p>The <i>st_blocks</i> field indicates the number of blocks allocated to the file, 512-byte units. (This may be smaller than <i>st_size/512</i> when the file has holes.)</p> <p>The <i>st_blksize</i> 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.)</p>	<p>stat(2)</p>
<p>SP-Miniklausur Manual-Auszug</p> <p>2019-05-16</p> <p>1</p>	<p>SP-Miniklausur Manual-Auszug</p> <p>2019-05-16</p> <p>1</p>	<p>SP-Miniklausur Manual-Auszug</p> <p>2019-05-16</p> <p>1</p>	<p>SP-Miniklausur Manual-Auszug</p> <p>2019-05-16</p> <p>1</p>

stat(2)

stat(2)

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)**, **fsstatat(2)**, **readlink(2)**, **utime(2)**, **capabilities(7)**, **symlink(7)**