

opendir/readdir(3)

opendir/readdir(3)

NAME

opendir – open a directory / readdir – read a directory

SYNOPSIS

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

```
#include <dirent.h>
```

```
DIR *opendir(const char *name);
```

```
struct dirent *readdir(DIR *dir);
```

```
int readdir_r(DIR *dirp, struct dirent *entry, struct dirent **result);
```

DESCRIPTION

The `opendir()` function opens a directory stream corresponding to the directory *name*, and returns a pointer to the directory stream. The stream is positioned at the first entry in the directory.

RETURN VALUE

The `opendir()` function returns a pointer to the directory stream or `NULL` if an error occurred.

DESCRIPTION readdir

The `readdir()` function returns a pointer to a dirent structure representing the next directory entry in the directory stream pointed to by *dir*. It returns `NULL` on reaching the end-of-file or if an error occurred.

DESCRIPTION readdir_r

The `readdir_r()` function initializes the structure referenced by *entry* and stores a pointer to this structure in *result*. On successful return, the pointer returned at *result* will have the same value as the argument *entry*. Upon reaching the end of the directory stream, this pointer will have the value `NULL`.

The data returned by `readdir()` is overwritten by subsequent calls to `readdir()` for the same directory stream.

The *dirent* structure is defined as follows:

```
struct dirent {
    long      d_ino;          /* inode number */
    off_t     d_off;         /* offset to the next dirent */
    unsigned short d_reclen; /* length of this record */
    unsigned char  d_type;   /* type of file */
    char         d_name[256]; /* filename */
};
```

RETURN VALUE

The `readdir()` function returns a pointer to a dirent structure, or `NULL` if an error occurs or end-of-file is reached.

`readdir_r()` returns 0 if successful or an error number to indicate failure.

ERRORS

EACCESS
Permission denied.

ENOENT

Directory does not exist, or *name* is an empty string.

ENOTDIR

name is not a directory.

fopen/fdopen/filenop(3)

fopen/fdopen/filenop(3)

NAME

fopen, fdopen, fileno – stream open functions

SYNOPSIS

```
#include <stdio.h>
```

```
FILE *fopen(const char *path, const char *mode);
```

```
FILE *fdopen(int fd, const char *mode);
```

```
int fileno(FILE *stream);
```

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 `flopen` function associates a stream with the existing file descriptor, *fd*. 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 *fd*, 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'd, and will be closed when the stream created by `flopen` is closed. The result of applying `flopen` to a shared memory object is undefined.

The function `fileno()` 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)`

FREAD(3)

FREAD(3)

malloc(3)

malloc(3)

NAME

fread, fwrite – binary stream input/output

SYNOPSIS

```
#include <stdio.h>
```

```
size_t fread(void *ptr, size_t size, size_t membh, FILE *stream);
```

```
size_t fwrite(const void *ptr, size_t size, size_t membh,
```

```
FILE *stream);
```

DESCRIPTION

The function `fread()` reads *membh* elements of data, each *size* bytes long, from the stream pointed to by *stream*, storing them at the location given by *ptr*.

The function `fwrite()` writes *membh* elements of data, each *size* bytes long, to the stream pointed to by *stream*, obtaining them from the location given by *ptr*.

For nonlocking counterparts, see `unlocked_stdio(3)`.

RETURN VALUE

`fread()` and `fwrite()` return the number of items successfully read or written (i.e., not the number of characters). If an error occurs, or the end-of-file is reached, the return value is a short item count (or zero).

`fread()` does not distinguish between end-of-file and error, and callers must use `feof(3)` and `ferror(3)` to determine which occurred.

CONFORMING TO

C89, POSIX.1-2001.

SEE ALSO

`read(2)`, `write(2)`, `feof(3)`, `ferror(3)`, `unlocked_stdio(3)`

NAME

calloc, malloc, free, realloc – Allocate and free dynamic memory

SYNOPSIS

```
#include <stdlib.h>
```

```
void *calloc(size_t membh, size_t size);
```

```
void *malloc(size_t size);
```

```
void free(void *ptr);
```

```
void *realloc(void *ptr, size_t size);
```

DESCRIPTION

`calloc()` allocates memory for an array of *membh* elements of *size* bytes each and returns a pointer to the allocated memory. The memory is set to zero.

`malloc()` allocates *size* bytes and returns a pointer to the allocated memory. The memory is not cleared.

`free()` frees the memory space pointed to by *ptr*, which must have been returned by a previous call to `malloc()`, `calloc()` or `realloc()`. Otherwise, or if `free(ptr)` has already been called before, undefined behaviour occurs. If *ptr* is `NULL`, no operation is performed.

`realloc()` changes the size of the memory block pointed to by *ptr* to *size* bytes. The contents will be unchanged to the minimum of the old and new sizes; newly allocated memory will be uninitialized. If *ptr* is `NULL`, the call is equivalent to `malloc(size)`; if *size* is equal to zero, the call is equivalent to `free(ptr)`. Unless *ptr* is `NULL`, it must have been returned by an earlier call to `malloc()`, `calloc()` or `realloc()`.

RETURN VALUE

For `calloc()` and `malloc()`, the value returned is a pointer to the allocated memory, which is suitably aligned for any kind of variable, or `NULL` if the request fails.

`free()` returns no value.

`realloc()` returns a pointer to the newly allocated memory, which is suitably aligned for any kind of variable and may be different from *ptr*, or `NULL` if the request fails. If *size* was equal to 0, either `NULL` or a pointer suitable to be passed to `free()` is returned. If `realloc()` fails the original block is left untouched - it is not freed or moved.

CONFORMING TO

ANSI-C

SEE ALSO

`brk(2)`, `posix_memalign(3)`

`pthread_create(pthread_t(3))`

`pthread_create(pthread_t(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_t(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 <XavierLeroy@inria.fr>

SEE ALSO

`pthread_join(3)`, `pthread_detach(3)`, `pthread_attr_t(3)`

`pthread_detach(3)`

`pthread_detach(3)`

NAME

`pthread_detach` – put a running thread in the detached state

SYNOPSIS

```
#include <pthread.h>
```

```
int pthread_detach(pthread_t th);
```

DESCRIPTION

`pthread_detach` put the thread *th* in the detached state. This guarantees that the memory resources consumed by *th* will be freed immediately when *th* terminates. However, this prevents other threads from synchronizing on the termination of *th* using `pthread_join`.

A thread can be created initially in the detached state, using the `detached` attribute to `pthread_create(3)`. In contrast, `pthread_detach` applies to threads created in the joinable state, and which need to be put in the detached state later.

After `pthread_detach` completes, subsequent attempts to perform `pthread_join` on *th* will fail. If another thread is already joining the thread *th* at the time `pthread_detach` is called, `pthread_detach` does nothing and leaves *th* in the joinable state.

RETURN VALUE

On success, 0 is returned. On error, a non-zero error code is returned.

ERRORS

ESRCH

No thread could be found corresponding to that specified by *th*

EINVAL

the thread *th* is already in the detached state

AUTHOR

Xavier Leroy <XavierLeroy@inria.fr>

SEE ALSO

`pthread_create(3)`, `pthread_join(3)`, `pthread_attr_setdetachstate(3)`

stat(2)

stat(2)

stat(2)

stat(2)

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 */
    blkcnt_t st_blksize; /* blocksize for the system I/O */
    btime_t st_btime; /* 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\(7\)](#)