

<p>opendir/readdir(3)</p> <p>NAME</p> <p>opendir — open a directory / readdir — read a directory</p> <p>SYNOPSIS</p> <pre>#include <sys/types.h> #include <dirent.h> DIR *opendir(const char *name); struct dirent *readdir(DIR *dir);</pre> <p>DESCRIPTION</p> <p>The <code>opendir()</code> 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</p> <p>The <code>opendir()</code> function returns a pointer to the directory stream or NULL if an error occurred.</p> <p>DESCRIPTION</p> <p>The <code>readdir()</code> 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 <code>readdir()</code> inside threads if the pointers passed as <i>dir</i> are created by distinct calls to <code>opendir()</code>. The data returned by <code>readdir()</code> is overwritten by subsequent calls to <code>readdir()</code> 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 */ off_t d_off; /* offset to the next dirent */ unsigned short d_reclen; /* length of this record */ unsigned char d_type; /* type of file; not supported by all filesystem types */ char d_name[256]; /* filename */ };</pre> <p>RETURN VALUE</p> <p>The <code>readdir()</code> function returns a pointer to a dirent structure, or NULL if an error occurs or end-of-file is reached.</p> <p>ERRORS</p> <p>EACCES Permission denied.</p> <p>ENOENT Directory does not exist, or <i>name</i> is an empty string.</p> <p>ENOTDIR <i>name</i> is not a directory.</p>	<p>opendir/readdir(3)</p> <p>NAME</p> <p>stat, fstat, lstat — get file status</p> <p>SYNOPSIS</p> <pre>#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);</pre> <p>Feature Test Macro Requirements for glibc (see <code>feature_test_macros(7)</code>):</p> <pre>lstat(_BSD_SOURCE _XOPEN_SOURCE >= 500</pre> <p>DESCRIPTION</p> <p>These functions return information about a file. No permissions are required on the file itself, but — in the case of <code>stat()</code> and <code>lstat()</code> — execute (search) permission is required on all of the directories in <i>path</i> that lead to the file.</p> <p><code>stat()</code> stats the file pointed to by <i>path</i> and fills in <i>buf</i>.</p> <p><code>lstat()</code> is identical to <code>stat()</code>, 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><code>fstat()</code> is identical to <code>stat()</code>, 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>stat(2)</p> <p>SP-Miniklausur Manual-Auszug</p> <p>2018-05-09</p> <p>1</p>
--	--	--

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\(7\)](#)