

opendir/readdir(3)

opendir/readdir(3)

dup(2)

dup(2)

**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 readir**

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**

**EACCES** Permission denied.

**ENOENT**

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

**ENOTDIR**

*name* is not a directory.

**NAME**

dup, dup2 – duplicate a file descriptor

**SYNOPSIS**

```
#include <unistd.h>
```

```
int dup(int oldfd);
```

```
int dup2(int oldfd, int newfd);
```

**DESCRIPTION**

`dup()` and `dup2()` create a copy of the file descriptor *oldfd*.

`dup()` uses the lowest-numbered unused descriptor for the new descriptor.

`dup2()` makes *newfd* be the copy of *oldfd*, closing *newfd* first if necessary, but note the following:

- \* If *oldfd* is not a valid file descriptor, then the call fails, and *newfd* is not closed.
- \* If *oldfd* is a valid file descriptor, and *newfd* has the same value as *oldfd*, then `dup2()` does nothing, and returns *newfd*.

After a successful return from `dup()` or `dup2()`, the old and new file descriptors may be used interchangeably. They refer to the same open file description (see `open(2)`) and thus share file offset and file status flags; for example, if the file offset is modified by using `seek(2)` on one of the descriptors, the offset is also changed for the other.

The two descriptors do not share file descriptor flags (the close-on-exec flag `(FD_CLOEXEC)`; see `fcntl(2)` for the duplicate descriptor is off.

**RETURN VALUE**

`dup()` and `dup2()` return the new descriptor, or `-1` if an error occurred (in which case, *errno* is set appropriately).

**ERRORS**

**EBADF**

*oldfd* isn't an open file descriptor, or *newfd* is out of the allowed range for file descriptors.

**EBUSY**

(Linux only) This may be returned by `dup2()` during a race condition with `open(2)` and `dup()`.

**EINTR**

The `dup2()` call was interrupted by a signal; see `signal(7)`.

**EMFILE**

The process already has the maximum number of file descriptors open and tried to open a new one.

**CONFORMING TO**

SV4, 4.3BSD, POSIX.1-2001.

**NOTES**

The error returned by `dup2()` is different from that returned by `fcntl(..., F_DUPED, ...)` when *newfd* is out of range. On some systems `dup2()` also sometimes returns `EINVAL` like `F_DUPED`.

If *newfd* was open, any errors that would have been reported at `close(2)` time are lost. A careful programmer will not use `dup2()` without closing *newfd* first.

**SEE ALSO**

`close(2)`, `fcntl(2)`, `open(2)`

**COLOPHON**

This page is part of release 3.05 of the Linux *man-pages* project. A description of the project, and information about reporting bugs, can be found at <http://www.kernel.org/doc/man-pages/>.

exec(2)

exec(2)

fileno(3)

fileno(3)

**NAME**

exec, execd, execv, execl, execlp, execlp - execute a file

**SYNOPSIS**

```
#include <unistd.h>

int execl(const char *path, const char *arg0, ..., const char *argn, char * /*NULL*/);
int execv(const char *path, char *const argv[]);
int execlp(const char *path, char *const argv[], ..., const char *argn,
char * /*NULL*/, char *const envp[]);
int execve(const char *path, char *const argv[], char *const envp[]);
int execlp(const char *file, const char *arg0, ..., const char *argn, char * /*NULL*/);
int execvp(const char *file, char *const argv[]);
```

**DESCRIPTION**

Each of the functions in the **exec** family overlays a new process image on an old process. The new process image is constructed from an ordinary, executable file. This file is either an executable object file, or a file of data for an interpreter. There can be no return from a successful call to one of these functions because the calling process image is overlaid by the new process image.

When a C program is executed, it is called as follows:

```
int main (int argc, char *argv[], char *envp[]);
```

where *argc* is the argument count, *argv* is an array of character pointers to the arguments themselves, and *envp* is an array of character pointers to the environment strings. As indicated, *argc* is at least one, and the first member of the array points to a string containing the name of the file.

The arguments *arg0*, ..., *argn* point to null-terminated character strings. These strings constitute the argument list available to the new process image. Conventionally at least *arg0* should be present. The *arg0* argument points to a string that is the same as *path* (or the last component of *path*). The list of argument strings is terminated by a **char \*0** argument.

The *argv* argument is an array of character pointers to null-terminated strings. These strings constitute the argument list available to the new process image. By convention, *argv* must have at least one member, and it should point to a string that is the same as *path* (or its last component). The *argv* argument is terminated by a null pointer.

The *path* argument points to a path name that identifies the new process file.

The *file* argument points to the new process file. If *file* does not contain a slash character, the path prefix for this file is obtained by a search of the directories passed in the **PATH** environment variable (see **environ(5)**).

File descriptors open in the calling process remain open in the new process.

Signals that are being caught by the calling process are set to the default disposition in the new process image (see **signal(3C)**). Otherwise, the new process image inherits the signal dispositions of the calling process.

**RETURN VALUES**

If a function in the **exec** family returns to the calling process, an error has occurred: the return value is **-1** and **errno** is set to indicate the error.

**NAME**

clearerr, feof, ferror, fileno - check and reset stream status

**SYNOPSIS**

```
#include <stdio.h>

void clearerr(FILE *stream);
int feof(FILE *stream);
int ferror(FILE *stream);
int fileno(FILE *stream);
```

**DESCRIPTION**

The function **clearerr()** clears the end-of-file and error indicators for the stream pointed to by *stream*.

The function **feof()** tests the end-of-file indicator for the stream pointed to by *stream*, returning non-zero if it is set. The end-of-file indicator can only be cleared by the function **clearerr()**.

The function **ferror()** tests the error indicator for the stream pointed to by *stream*, returning non-zero if it is set. The error indicator can only be reset by the **clearerr()** function.

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

For non-locking counterparts, see **unlocked\_stdio(3)**.

**ERRORS**

These functions should not fail and do not set the external variable *errno*. (However, in case **fileno()** detects that its argument is not a valid stream, it must return **-1** and set *errno* to **EBADF**.)

**CONFORMING TO**

The functions **clearerr()**, **feof()**, and **ferror()** conform to C89 and C99.

**SEE ALSO**

**open(2)**, **fdopen(3)**, **stdio(3)**, **unlocked\_stdio(3)**

fopen/fdopen(3)

fopen/fdopen(3)

sigaction(2)

sigaction(2)

**NAME**

fopen, fdopen – stream open functions

**SYNOPSIS**

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

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

**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, *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'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.

**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)

**NAME**

sigaction – POSIX signal handling functions.

**SYNOPSIS**

```
#include <signal.h>
```

```
int sigaction(int signum, 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.

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

If *act* is non-null, the new action for signal *signum* 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);  
    void (*sa_sigaction)(int, siginfo_t *, void *);  
    sigset_t sa_mask;  
    int sa_flags;  
    void (*sa_restorer)(void);  
};
```

On some architectures a union is involved – do not assign to both *sa\_handler* and *sa\_sigaction*.

The *sa\_restorer* element is obsolete and should not be used. POSIX does not specify a *sa\_restorer* element.

*sa\_handler* specifies the action to be associated with *signum* 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 *signum* 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.

**RETURN VALUES**

**sigaction** returns 0 on success and -1 on 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).

sigsuspend/sigprocmask(2)

sigsuspend/sigprocmask(2)

#### NAME

sigprocmask – change and/or examine caller's signal mask  
sigsuspend – install a signal mask and suspend caller until signal

#### SYNOPSIS

```
#include <signal.h>
int sigprocmask(int how, const sigset_t *set, sigset_t *oset);
int sigsuspend(const sigset_t *set);
```

#### DESCRIPTION

The **sigprocmask()** function is used to examine and/or change the caller's signal mask. If the value is **SIG\_BLOCK**, the set pointed to by the argument *set* is added to the current signal mask. If the value is **SIG\_UNBLOCK**, the set pointed by the argument *set* is removed from the current signal mask. If the value is **SIG\_SETMASK**, the current signal mask is replaced by the set pointed to by the argument *set*. If the argument *oset* is not NULL, the previous mask is stored in the space pointed to by *oset*. If the value of the argument *set* is NULL, the value *how* is not significant and the caller's signal mask is unchanged; thus, the call can be used to inquire about currently blocked signals.

If there are any pending unblocked signals after the call to **sigprocmask()**, at least one of those signals will be delivered before the call to **sigprocmask()** returns.

It is not possible to block those signals that cannot be ignored this restriction is silently imposed by the system. See **sigaction(2)**.

If **sigprocmask()** fails, the caller's signal mask is not changed.

#### RETURN VALUES

On success, **sigprocmask()** returns **0**. On failure, it returns **-1** and sets **errno** to indicate the error.

#### ERRORS

**sigprocmask()** fails if any of the following is true:

**EFAULT** *set* or *oset* points to an illegal address.

**EINVAL** The value of the *how* argument is not equal to one of the defined values.

#### DESCRIPTION

**sigsuspend()** replaces the caller's signal mask with the set of signals pointed to by the argument *set* and then suspends the caller until delivery of a signal whose action is either to execute a signal catching function or to terminate the process.

If the action is to terminate the process, **sigsuspend()** does not return. If the action is to execute a signal catching function, **sigsuspend()** returns after the signal catching function returns. On return, the signal mask is restored to the set that existed before the call to **sigsuspend()**.

It is not possible to block those signals that cannot be ignored (see **signal(5)**); this restriction is silently imposed by the system.

#### RETURN VALUES

Since **sigsuspend()** suspends process execution indefinitely, there is no successful completion return value. On failure, it returns **-1** and sets **errno** to indicate the error.

#### ERRORS

**sigsuspend()** fails if either of the following is true:

**EFAULT** *set* points to an illegal address.

**EINTR** A signal is caught by the calling process and control is returned from the signal catching function.

#### SEE ALSO

**sigaction(2)**, **sigsetops(3C)**,

sigsetops(3C)

sigsetops(3C)

#### NAME

sigsetops, sigemptyset, sigfillset, sigaddset, sigdelset, sigismember – manipulate sets of signals

#### SYNOPSIS

```
#include <signal.h>
int sigemptyset(sigset_t *set);
int sigfillset(sigset_t *set);
int sigaddset(sigset_t *set, int signo);
int sigdelset(sigset_t *set, int signo);
int sigismember(sigset_t *set, int signo);
```

#### DESCRIPTION

These functions manipulate *sigset\_t* data types, representing the set of signals supported by the implementation.

**sigemptyset()** initializes the set pointed to by *set* to exclude all signals defined by the system.

**sigfillset()** initializes the set pointed to by *set* to include all signals defined by the system.

**sigaddset()** adds the individual signal specified by the value of *signo* to the set pointed to by *set*.

**sigdelset()** deletes the individual signal specified by the value of *signo* from the set pointed to by *set*.

**sigismember()** checks whether the signal specified by the value of *signo* is a member of the set pointed to by *set*.

Any object of type *sigset\_t* must be initialized by applying either **sigemptyset()** or **sigfillset()** before applying any other operation.

#### RETURN VALUES

Upon successful completion, the **sigismember()** function returns a value of one if the specified signal is a member of the specified set, or a value of 0 if it is not. Upon successful completion, the other functions return a value of 0. Otherwise a value of **-1** is returned and **errno** is set to indicate the error.

#### ERRORS

**sigaddset()**, **sigdelset()**, and **sigismember()** will fail if the following is true:

**EINVAL** The value of the *signo* argument is not a valid signal number.

**sigfillset()** will fail if the following is true:

**EFAULT** The *set* argument specifies an invalid address.

#### SEE ALSO

**sigaction(2)**, **sigpending(2)**, **sigprocmask(2)**, **sigsuspend(2)**, **attributes(5)**, **signal(5)**

printf(3)

printf(3)

NAME

printf, fprintf, sprintf, snprintf, vprintf, vsprintf, vsnprintf – formatted output conversion

SYNOPSIS

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

```
int printf(const char *format, ...);
int fprintf(FILE *stream, const char *format, ...);
int sprintf(char *str, const char *format, ...);
int snprintf(char *str, size_t size, const char *format, ...);
```

...

DESCRIPTION

The functions in the **printf()** family produce output according to a *format* as described below. The functions **printf()** and **vprintf()** write output to *stdout*, the standard output stream; **fprintf()** and **vfprintf()** write output to the given output *stream*; **sprintf()**, **snprintf()**, **vsprintf()** and **vsnprintf()** write to the character string *str*.

The functions **sprintf()** and **vsnprintf()** write at most *size* bytes (including the trailing null byte '\0') to *str*.

These eight functions write the output under the control of a *format* string that specifies how subsequent arguments (or arguments accessed via the variable-length argument facilities of **stdarg(3)**) are converted for output.

**Return value**

Upon successful return, these functions return the number of characters printed (not including the trailing '\0' used to end output to strings).

The functions **sprintf()** and **vsnprintf()** do not write more than *size* bytes (including the trailing '\0'). If the output was truncated due to this limit then the return value is the number of characters (not including the trailing '\0') which would have been written to the final string if enough space had been available. Thus, a return value of *size* or more means that the output was truncated. (See also below under NOTES.)

If an output error is encountered, a negative value is returned.

**Format of the format string**

The format string is a character string, beginning and ending in its initial shift state, if any. The format string is composed of zero or more directives: ordinary characters (not %), which are copied unchanged to the output stream, and conversion specifications, each of which results in fetching zero or more subsequent arguments. Each conversion specification is introduced by the character %, and ends with a *conversion specifier*. In between there may be (in this order) zero or more *flags*, an optional minimum *field width*, an optional *precision* and an optional *length modifier*.

**The conversion specifier**

A character that specifies the type of conversion to be applied. An example for a conversion specifier is:

**s**      The *convst char \** argument is expected to be a pointer to an array of character type (pointer to a string). Characters from the array are written up to (but not including) a terminating null byte ('\0'); if a precision is specified, no more than the number specified are written. If a precision is given, no null byte need be present; if the precision is not specified, or is greater than the size of the array, the array must contain a terminating null byte.

SEE ALSO

**printf(1)**, **asprintf(3)**, **dprintf(3)**, **scanf(3)**, **setlocale(3)**, **wcrtomb(3)**, **wprintf(3)**, **locale(5)**

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 */
    blksize_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-write.)

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 "rountine" 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\(7\)](#)

waitpid(2)

waitpid(2)

#### NAME

waitpid — wait for child process to change state

#### SYNOPSIS

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

```
pid_t waitpid(pid_t pid, int *stat_loc, int options);
```

#### DESCRIPTION

**waitpid()** suspends the calling process until one of its children changes state; if a child process changed state prior to the call to **waitpid()**, return is immediate. *pid* specifies a set of child processes for which status is requested.

If *pid* is equal to **(pid\_t)-1**, status is requested for any child process.

If *pid* is greater than **(pid\_t)0**, it specifies the process ID of the child process for which status is requested.

If *pid* is equal to **(pid\_t)0** status is requested for any child process whose process group ID is equal to that of the calling process.

If *pid* is less than **(pid\_t)-1**, status is requested for any child process whose process group ID is equal to the absolute value of *pid*.

If **waitpid()** returns because the status of a child process is available, then that status may be evaluated with the macros defined by [wstat\(5\)](#). If the calling process had specified a non-zero value of *stat\_loc*, the status of the child process will be stored in the location pointed to by *stat\_loc*.

The *options* argument is constructed from the bitwise inclusive OR of zero or more of the following flags, defined in the header `<sys/wait.h>`:

##### WCONTINUED

The status of any continued child process specified by *pid*, whose status has not been reported since it continued, is also reported to the calling process.

##### WNOHANG

**waitpid()** will not suspend execution of the calling process if status is not immediately available for one of the child processes specified by *pid*.

##### WNOVAIT

Keep the process whose status is returned in *stat\_loc* in a waitable state. The process may be waited for again with identical results.

#### RETURN VALUES

If **waitpid()** returns because the status of a child process is available, this function returns a value equal to the process ID of the child process for which status is reported. If **waitpid()** returns due to the delivery of a signal to the calling process, `-1` is returned and *errno* is set to **EINTR**. If this function was invoked with **WNOHANG** set in *options*, it has at least one child process specified by *pid* for which status is not available, and status is not available for any process specified by *pid*, `0` is returned. Otherwise, `-1` is returned, and *errno* is set to indicate the error.

#### ERRORS

**waitpid()** will fail if one or more of the following is true:

##### ECHILD

The process or process group specified by *pid* does not exist or is not a child of the calling process or can never be in the states specified by *options*.

##### EINTR

**waitpid()** was interrupted due to the receipt of a signal sent by the calling process.

##### EINVAL

An invalid value was specified for *options*.

#### SEE ALSO

[exec\(2\)](#), [exit\(2\)](#), [fork\(2\)](#), [sigaction\(2\)](#), [wstat\(5\)](#)