

connect(2)

connect(2)

NAME

connect – initiate a connection on a socket

SYNOPSIS

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

```
int connect(int sockfd, const struct sockaddr *serv_addr, socklen_t addrlen);
```

DESCRIPTION

The file descriptor *sockfd* must refer to a socket. If the socket is of type **SOCK_DGRAM** then the *serv_addr* address is the address to which datagrams are sent by default, and the only address from which datagrams are received. If the socket is of type **SOCK_STREAM** or **SOCK_SEQPACKET**, this call attempts to make a connection to another socket. The other socket is specified by *serv_addr*, which is an address (of length *addrlen*) in the communications space of the socket. Each communications space interprets the *serv_addr* parameter in its own way.

Generally, connection-based protocol sockets may successfully **connect** only once; connectionless protocol sockets may use **connect** multiple times to change their association. Connectionless sockets may dissolve the association by connecting to an address with the *sa_family* member of **sockaddr** set to **AF_UNSPEC**.

RETURN VALUE

If the connection or binding succeeds, zero is returned. On error, -1 is returned, and *errno* is set appropriately.

ERRORS

The following are general socket errors only. There may be other domain-specific error codes.

EBADF

The file descriptor is not a valid index in the descriptor table.

EFAULT

The socket structure address is outside the user's address space.

ENOTSOCK

The file descriptor is not associated with a socket.

EISCONN

The socket is already connected.

ECONNREFUSED

No one listening on the remote address.

ENETUNREACH

Network is unreachable.

EADDRINUSE

Local address is already in use.

EAFNOSUPPORT

The passed address didn't have the correct address family in its *sa_family* field.

EACCES, EPERM

The user tried to connect to a broadcast address without having the socket broadcast flag enabled or the connection request failed because of a local firewall rule.

SEE ALSO

accept(2), **bind(2)**, **listen(2)**, **socket(2)**, **getsockname(2)**

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);
```

DESCRIPTION opendir

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.

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.

ERRORS

EACCES

Permission denied.

EMFILE

Too many file descriptors in use by process.

ENFILE

Too many files are currently open in the system.

ENOENT

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

ENOMEM

Insufficient memory to complete the operation.

ENOTDIR

name is not a directory.

SEE ALSO

open(2), **readdir(3)**, **closedir(3)**, **rewinddir(3)**, **seekdir(3)**, **telldir(3)**, **scandir(3)**

NAME

fdopen – associate a stream with a file descriptor

SYNOPSIS

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

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

DESCRIPTION

The **fdopen()** function associates a stream with a file descriptor *fildes*, 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 *fildes*.

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 <i>fildes</i> argument is not a valid file descriptor.
EINVAL	The <i>mode</i> 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)**

NAME

gets, fgets – get a string from a stream
fputs, puts – output of strings

SYNOPSIS

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

```
char *gets(char *s);
```

```
char *fgets(char *s, int n, FILE *stream);
```

```
int fputs(const char *s, FILE *stream);
```

```
int puts(const char *s);
```

DESCRIPTION gets/fgets

The **gets()** function reads characters from the standard input stream (see **intro(3)**), **stdin**, into the array pointed to by *s*, until a newline character is read or an end-of-file condition is encountered. The newline character is discarded and the string is terminated with a null character.

The **fgets()** function reads characters from the *stream* into the array pointed to by *s*, until *n*-1 characters are read, or a newline character is read and transferred to *s*, or an end-of-file condition is encountered. The string is then terminated with a null character.

When using **gets()**, if the length of an input line exceeds the size of *s*, indeterminate behavior may result. For this reason, it is strongly recommended that **gets()** be avoided in favor of **fgets()**.

RETURN VALUES

If end-of-file is encountered and no characters have been read, no characters are transferred to *s* and a null pointer is returned. If a read error occurs, such as trying to use these functions on a file that has not been opened for reading, a null pointer is returned and the error indicator for the stream is set. If end-of-file is encountered, the **EOF** indicator for the stream is set. Otherwise *s* is returned.

ERRORS

The **gets()** and **fgets()** functions will fail if data needs to be read and:

EIO	The file is a regular file and an attempt was made to read at or beyond the offset maximum associated with the corresponding <i>stream</i> .
------------	--

DESCRIPTION puts/fputs

fputs() writes the string *s* to *stream*, without its trailing **^M**.

puts() writes the string *s* and a trailing newline to *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

puts() and **fputs()** return a non - negative number on success, or **EOF** on error.

NAME

ip – Linux IPv4 protocol implementation

SYNOPSIS

```
#include <sys/socket.h>
#include <netinet/in.h>
```

```
tcp_socket = socket(PF_INET, SOCK_STREAM, 0);
raw_socket = socket(PF_INET, SOCK_RAW, protocol);
udp_socket = socket(PF_INET, SOCK_DGRAM, protocol);
```

DESCRIPTION

The programmer's interface is BSD sockets compatible. For more information on sockets, see [socket\(7\)](#).

An IP socket is created by calling the [socket\(2\)](#) function as `socket(PF_INET, socket_type, protocol)`. Valid socket types are `SOCK_STREAM` to open a [tcp\(7\)](#) socket, `SOCK_DGRAM` to open a [udp\(7\)](#) socket, or `SOCK_RAW` to open a [raw\(7\)](#) socket to access the IP protocol directly. *protocol* is the IP protocol in the IP header to be received or sent. The only valid values for *protocol* are `0` and `IPPROTO_TCP` for TCP sockets and `0` and `IPPROTO_UDP` for UDP sockets.

When a process wants to receive new incoming packets or connections, it should bind a socket to a local interface address using [bind\(2\)](#). Only one IP socket may be bound to any given local (address, port) pair. When `INADDR_ANY` is specified in the bind call the socket will be bound to *all* local interfaces. When [listen\(2\)](#) or [connect\(2\)](#) are called on a unbound socket the socket is automatically bound to a random free port with the local address set to `INADDR_ANY`.

ADDRESS FORMAT

An IP socket address is defined as a combination of an IP interface address and a port number. The basic IP protocol does not supply port numbers, they are implemented by higher level protocols like [tcp\(7\)](#).

```
struct sockaddr_in {
    sa_family_t    sin_family; /* address family: AF_INET */
    u_int16_t      sin_port;   /* port in network byte order */
    struct in_addr sin_addr;   /* internet address */
};
/* Internet address. */
struct in_addr {
    u_int32_t      s_addr;    /* address in network byte order */
};
```

sin_family is always set to `AF_INET`. This is required; in Linux 2.2 most networking functions return `EINVAL` when this setting is missing. *sin_port* contains the port in network byte order. The port numbers below 1024 are called *reserved ports*. Only processes with effective user id 0 or the `CAP_NET_BIND_SERVICE` capability may [bind\(2\)](#) to these sockets.

sin_addr is the IP host address. The *addr* member of `struct in_addr` contains the host interface address in network order. `in_addr` should be only accessed using the [inet_aton\(3\)](#), [inet_addr\(3\)](#), [inet_makeaddr\(3\)](#) library functions or directly with the name resolver (see [gethostbyname\(3\)](#)).

Note that the address and the port are always stored in network order. In particular, this means that you need to call [htons\(3\)](#) on the number that is assigned to a port. All address/port manipulation functions in the standard library work in network order.

SEE ALSO

[sendmsg\(2\)](#), [recvmsg\(2\)](#), [socket\(7\)](#), [netlink\(7\)](#), [tcp\(7\)](#), [udp\(7\)](#), [raw\(7\)](#), [ipfw\(7\)](#)

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\(3\)](#),

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 sigprocmask

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

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)**,

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)**

NAME

socket – create an endpoint for communication

SYNOPSIS

```
cc [ flag ... ] file ... -lsocket -lnsl [ library ... ]
```

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

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

```
int socket(int domain, int type, int protocol);
```

DESCRIPTION

socket() creates an endpoint for communication and returns a descriptor.

The *domain* parameter specifies a communications domain within which communication will take place; this selects the protocol family which should be used. The protocol family generally is the same as the address family for the addresses supplied in later operations on the socket. These families are defined in the include file `<sys/socket.h>`. There must be an entry in the `netconfig(4)` file for at least each protocol family and type required. If *protocol* has been specified, but no exact match for the tuple family, type, protocol is found, then the first entry containing the specified family and type with zero for protocol will be used. The currently understood formats are:

```
PF_UNIX   UNIX system internal protocols
```

```
PF_INET   ARPA Internet protocols
```

The socket has the indicated *type*, which specifies the communication semantics. Currently defined types are:

```
SOCK_STREAM
```

```
SOCK_DGRAM
```

```
SOCK_RAW
```

```
SOCK_SEQPACKET
```

```
SOCK_RDM
```

A **SOCK_STREAM** type provides sequenced, reliable, two-way connection-based byte streams. An out-of-band data transmission mechanism may be supported. A **SOCK_DGRAM** socket supports datagrams (connectionless, unreliable messages of a fixed (typically small) maximum length). A **SOCK_SEQPACKET** socket may provide a sequenced, reliable, two-way connection-based data transmission path for datagrams of fixed maximum length; a consumer may be required to read an entire packet with each read system call. This facility is protocol specific, and presently not implemented for any protocol family. **SOCK_RAW** sockets provide access to internal network interfaces. The types **SOCK_RAW**, which is available only to the super-user, and **SOCK_RDM**, for which no implementation currently exists, are not described here.

protocol specifies a particular protocol to be used with the socket. Normally only a single protocol exists to support a particular socket type within a given protocol family. However, multiple protocols may exist, in which case a particular protocol must be specified in this manner. The protocol number to use is particular to the “communication domain” in which communication is to take place. If a protocol is specified by the caller, then it will be packaged into a socket level option request and sent to the underlying protocol layers.

Sockets of type **SOCK_STREAM** are full-duplex byte streams, similar to pipes. A stream socket must be in a *connected* state before any data may be sent or received on it. A connection to another socket is created with a `connect(3N)` call. Once connected, data may be transferred using `read(2)` and `write(2)` calls or some variant of the `send(3N)` and `recv(3N)` calls. When a session has been completed, a `close(2)` may be performed. Out-of-band data may also be transmitted as described on the `send(3N)` manual page and received as described on the `recv(3N)` manual page.

The communications protocols used to implement a **SOCK_STREAM** insure that data is not lost or duplicated. If a piece of data for which the peer protocol has buffer space cannot be successfully transmitted within a reasonable length of time, then the connection is considered broken and calls will indicate an error with `-1` returns and with `ETIMEDOUT` as the specific code in the global variable `errno`. The protocols optionally keep sockets “warm” by forcing transmissions roughly every minute in the absence of other

activity. An error is then indicated if no response can be elicited on an otherwise idle connection for a extended period (for instance 5 minutes). A `SIGPIPE` signal is raised if a process sends on a broken stream; this causes naive processes, which do not handle the signal, to exit.

SOCK_SEQPACKET sockets employ the same system calls as **SOCK_STREAM** sockets. The only difference is that `read(2)` calls will return only the amount of data requested, and any remaining in the arriving packet will be discarded.

SOCK_DGRAM and **SOCK_RAW** sockets allow datagrams to be sent to correspondents named in `sendto(3N)` calls. Datagrams are generally received with `recvfrom(3N)`, which returns the next datagram with its return address.

An `fcntl(2)` call can be used to specify a process group to receive a `SIGURG` signal when the out-of-band data arrives. It may also enable non-blocking I/O and asynchronous notification of I/O events with `SIGIO` signals.

The operation of sockets is controlled by socket level *options*. These options are defined in the file `<sys/socket.h>`. `setsockopt(3N)` and `getsockopt(3N)` are used to set and get options, respectively.

RETURN VALUES

A `-1` is returned if an error occurs. Otherwise the return value is a descriptor referencing the socket.

ERRORS

The `socket()` call fails if:

EACCESS	Permission to create a socket of the specified type and/or protocol is denied.
EMFILE	The per-process descriptor table is full.
ENOMEM	Insufficient user memory is available.
ENOSR	There were insufficient STREAMS resources available to complete the operation.
EPROTONOSUPPORT	The protocol type or the specified protocol is not supported within this domain.

SEE ALSO

`close(2)`, `fcntl(2)`, `ioctl(2)`, `read(2)`, `write(2)`, `accept(3N)`, `bind(3N)`, `connect(3N)`, `getsockname(3N)`, `getsockopt(3N)`, `listen(3N)`, `recv(3N)`, `setsockopt(3N)`, `send(3N)`, `shutdown(3N)`, `socketpair(3N)`, `attributes(5)`, `in(5)`, `socket(5)`

unlink(2)

unlink(2)

NAME

unlink – remove directory entry

SYNOPSIS

```
#include <unistd.h>
int unlink(const char *path);
```

DESCRIPTION

The **unlink()** function removes a link to a file. It removes the link named by the pathname pointed to by *path* and decrements the link count of the file referenced by the link.

When the file's link count becomes 0 and no process has the file open, the space occupied by the file will be freed and the file will no longer be accessible. If one or more processes have the file open when the last link is removed, the link will be removed before **unlink()** returns, but the removal of the file contents will be postponed until all references to the file are closed.

RETURN VALUES

Upon successful completion, **0** is returned. Otherwise, **-1** is returned and **errno** is set to indicate the error.

ERRORS

The **unlink()** function will fail and not unlink the file if:

EACCES	Search permission is denied for a component of the <i>path</i> prefix.
EACCES	Write permission is denied on the directory containing the link to be removed.
ENOENT	The named file does not exist or is a null pathname.
ENOTDIR	A component of the <i>path</i> prefix is not a directory.
EPERM	The named file is a directory and the effective user of the calling process is not super-user.

SEE ALSO

rm(1), **close(2)**, **link(2)**, **open(2)**, **rmdir(2)**,

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 <i>pid</i> , 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 <i>pid</i> .
WNOWAIT	Keep the process whose status is returned in <i>stat_loc</i> 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 <i>pid</i> does not exist or is not a child of the calling process or can never be in the states specified by <i>options</i> .
EINTR	waitpid() was interrupted due to the receipt of a signal sent by the calling process.
EINVAL	An invalid value was specified for <i>options</i> .

SEE ALSO

exec(2), **exit(2)**, **fork(2)**, **sigaction(2)**, **wstat(5)**