

accept(2)

accept(2)

NAME
accept – accept a connection on a socket

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

`int accept(int s, struct sockaddr *addr, int *addrlen);`

DESCRIPTION

The argument *s* is a socket that has been created with `socket(3N)` and bound to an address with `bind(3N)`, and that is listening for connections after a call to `listen(3N)`. The `accept()` function extracts the first connection on the queue of pending connections, creates a new socket with the properties of *s*, and allocates a new file descriptor, *ns*, for the socket. If no pending connections are present on the queue and the socket is not marked as non-blocking, `accept()` blocks the caller until a connection is present. If the socket is marked as non-blocking and no pending connections are present on the queue, `accept()` returns an error as described below. The `accept()` function uses the `netconfig(4)` file to determine the STREAMS device file name associated with *s*. This is the device on which the connect indication will be accepted. The accepted socket, *ns*, is used to read and write data to and from the socket that connected to *s*; it is not used to accept more connections. The original socket (*s*) remains open for accepting further connections.

The argument *addr* is a result parameter that is filled in with the address of the connecting entity as it is known to the communications layer. The exact format of the *addr* parameter is determined by the domain in which the communication occurs.

The argument *addrlen* is a value-result parameter. Initially, it contains the amount of space pointed to by *addr*; on return it contains the length in bytes of the address returned.

The `accept()` function is used with connection-based socket types, currently with `SOCK_STREAM`.

It is possible to `select(3C)` or `poll(2)` a socket for the purpose of an `accept()` by selecting or polling it for a read. However, this will only indicate when a connect indication is pending; it is still necessary to call `accept()`.

RETURN VALUES

The `accept()` function returns `-1` on error. If it succeeds, it returns a non-negative integer that is a descriptor for the accepted socket.

ERRORS

`accept()` will fail if:

EBADF The descriptor is invalid.

EINTR The accept attempt was interrupted by the delivery of a signal.

EMFILE The per-process descriptor table is full.

ENODEV The protocol family and type corresponding to *s* could not be found in the `netconfig` file file.

ENOMEM There was insufficient user memory available to complete the operation.

EPROTO A protocol error has occurred; for example, the STREAMS protocol stack has not been initialized or the connection has already been released.

EWOLDBLOCK The socket is marked as non-blocking and no connections are present to be accepted.

SEE ALSO

`poll(2)`, `bind(3N)`, `connect(3N)`, `listen(3N)`, `select(3C)`, `socket(3N)`, `netconfig(4)`, `attributes(5)`, `socket(5)`

bind(2)

bind(2)

NAME
bind – bind a name to a socket

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

`int bind(int s, const struct sockaddr *name, int namelen);`

DESCRIPTION

`bind()` assigns a name to an unnamed socket. When a socket is created with `socket(3N)`, it exists in a name space (address family) but has no name assigned. `bind()` requests that the name pointed to by *name* be assigned to the socket.

RETURN VALUES

If the bind is successful, `0` is returned. A return value of `-1` indicates an error, which is further specified in the global `errno`.

ERRORS

The `bind()` call will fail if:

EACCES The requested address is protected and the current user has inadequate permission to access it.

EADDRINUSE The specified address is already in use.

EADDRNOTAVAIL The specified address is not available on the local machine.

EBADF *s* is not a valid descriptor.

EINVAL *namelen* is not the size of a valid address for the specified address family.

EINVAL The socket is already bound to an address.

ENOSR There were insufficient STREAMS resources for the operation to complete.

ENOTSOCK *s* is a descriptor for a file, not a socket.

The following errors are specific to binding names in the UNIX domain:

EACCES Search permission is denied for a component of the path prefix of the pathname in *name*.

EIO An I/O error occurred while making the directory entry or allocating the inode.

EISDIR A null pathname was specified.

ELOOP Too many symbolic links were encountered in translating the pathname in *name*.

ENOENT A component of the path prefix of the pathname in *name* does not exist.

ENOTDIR A component of the path prefix of the pathname in *name* is not a directory.

EROFS The inode would reside on a read-only file system.

SEE ALSO

`unlink(2)`, `socket(3N)`, `attributes(5)`, `socket(5)`

NOTES

Binding a name in the UNIX domain creates a socket in the file system that must be deleted by the caller when it is no longer needed (using `unlink(2)`).

The rules used in name binding vary between communication domains.

dup(2)

dup(2)

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 lseek(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). The close-on-exec flag (FD_CLOEXEC; seefcntl(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

SVr4, 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)

NAME

exec, execd, execv, execl, execlp, execlp, 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[]);
int execlp(const char *path, char *const argv[]);
int execlp(const char *file, const char *arg0, ..., const char *argn, char *#NULL*);
int execlp(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 *argv*0, ..., *argv*n point to null-terminated character strings. These strings constitute the argument list available to the new process image. Conventionally at least *argv*0 should be present. The *argv*0 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*n 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*n 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*n 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.

fopen/fdopen/filenop(3)

fopen/fdopen/filenop(3)

gets(3)

gets(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'ed, 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 **open** 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

gets, fgets – get a string from a stream

puts, 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:

EOVERFLOW The file is a regular file and an attempt was made to read at or beyond the offset maximum associated with the corresponding *stream*.

DESCRIPTION

fputs() writes the string *s* to *stream*, without its trailing '\0'.

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.

socket(2) / ipv6(7)

socket(2) / ipv6(7)

listen(2)

listen(2)

NAME

ipv6, PF_INET6 – Linux IPv6 protocol implementation

SYNOPSIS

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

```
tcp6_socket = socket(PF_INET6, SOCK_STREAM, 0);
raw6_socket = socket(PF_INET6, SOCK_RAW, protocol);
udp6_socket = socket(PF_INET6, SOCK_DGRAM, protocol);
```

DESCRIPTION

Linux 2.2 optionally implements the Internet Protocol version 6. This man page contains a description of the IPv6 basic API as implemented by the Linux kernel and glibc 2.1. The interface is based on the BSD sockets interface; see [socket\(7\)](#).

The IPv6 API aims to be mostly compatible with the [ip\(7\)](#) v4 API. Only differences are described in this man page.

To bind an **AF_INET6** socket to any process the local address should be copied from the *inaddr_any* variable which has *in6_addr* type. In static initializations **IN6ADDR_ANY_INIT** may also be used, which expands to a constant expression. Both of them are in network order.

The IPv6 loopback address (:::1) is available in the global *inaddr_loopback* variable. For initializations **IN6ADDR_LOOPBACK_INIT** should be used.

IPv4 connections can be handled with the v6 API by using the v4-mapped-on-v6 address type; thus a program only needs only to support this API type to support both protocols. This is handled transparently by the address handling functions in libc.

IPv4 and IPv6 share the local port space. When you get an IPv4 connection or packet to a IPv6 socket its source address will be mapped to v6 and it will be mapped to v6.

Address Format

```
struct sockadd_in6 {
    uint16_t    sin6_family; /* AF_INET6 */
    uint16_t    sin6_port; /* port number */
    uint32_t    sin6_flowinfo; /* IPv6 flow information */
    struct in6_addr sin6_addr; /* IPv6 address */
    uint32_t    sin6_scope_id; /* Scope ID (new in 2.4) */
};

struct in6_addr {
    unsigned char s6_addr[16]; /* IPv6 address */
};
```

sin6_family is always set to **AF_INET6**; *sin6_port* is the protocol port (see [sin_port](#) in [ip\(7\)](#)); *sin6_flowinfo* is the IPv6 flow identifier; *sin6_addr* is the 128-bit IPv6 address. *sin6_scope_id* is an ID of depending of on the scope of the address. It is new in Linux 2.4. Linux only supports it for link scope addresses, in that case *sin6_scope_id* contains the interface index (see [netdevice\(7\)](#))

NOTES

The *sockadd_in6* structure is bigger than the generic *sockaddr*. Programs that assume that all address types can be stored safely in a *struct sockaddr* need to use *struct sockaddr_storage* for that instead.

SEE ALSO

[cmsnsg\(3\)](#), [ip\(7\)](#)

NAME

listen – listen for connections on a socket

SYNOPSIS

```
#include <sys/types.h> /* See NOTES */
#include <sys/socket.h>
```

```
int listen(int sockfd, int backlog);
```

DESCRIPTION

[listen\(\)](#) marks the socket referred to by *sockfd* as a passive socket, that is, as a socket that will be used to accept incoming connection requests using [accept\(2\)](#).

The *sockfd* argument is a file descriptor that refers to a socket of type **SOCK_STREAM** or **SOCK_SEQPACKET**.

The *backlog* argument defines the maximum length to which the queue of pending connections for *sockfd* may grow. If a connection request arrives when the queue is full, the client may receive an error with an indication of **ECONNREFUSED** or, if the underlying protocol supports retransmission, the request may be ignored so that a later reattempt at connection succeeds.

RETURN VALUE

On success, zero is returned. On error, `-1` is returned, and *errno* is set appropriately.

ERRORS

EADDRINUSE

Another socket is already listening on the same port.

EBADF

The argument *sockfd* is not a valid descriptor.

ENOTSOCK

The argument *sockfd* is not a socket.

NOTES

To accept connections, the following steps are performed:

1. A socket is created with [socket\(2\)](#).
2. The socket is bound to a local address using [bind\(2\)](#), so that other sockets may be [connect\(2\)](#)ed to it.
3. A willingness to accept incoming connections and a queue limit for incoming connections are specified with [listen\(\)](#).
4. Connections are accepted with [accept\(2\)](#).

If the *backlog* argument is greater than the value in */proc/sys/net/core/somaxconn*, then it is silently truncated to that value; the default value in this file is 128.

EXAMPLE

See [bind\(2\)](#).

SEE ALSO

[accept\(2\)](#), [bind\(2\)](#), [connect\(2\)](#), [socket\(2\)](#), [socket\(7\)](#)

sigaction(2)

sigaction(2)

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`, `SIGTIN` 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

`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)`.

sigsetops(3C)

sigsetops(3C)

waitpid(2)

waitpid(2)

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:

EINVAL The *set* argument specifies an invalid address.

SEE ALSO

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

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_t0**, 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*.

WNOWAIT

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)**