

accept(3)

accept(3)

bind(3)

bind(3)

NAME

accept – accept a connection on a socket

NAME

bind – bind a name to a socket

SYNOPSIS

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

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

SYNOPSIS

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

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

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 *ns*; it is not used to accept more connections. The original socket (*s*) remains open for accepting further connections.

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.

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.

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**.

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.

ERRORS

The **bind()** call will fail if:

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

- 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.

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()**.

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

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.

- 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.

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

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 differ between communication domains.

SEE ALSO

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

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**; 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

SVr4, 4.3BSD, POSIX.1-2001.

NOTES

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

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/>.

NAME

exec, execl, execv, execl, execve, execlp, execvp – 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 execl(const char *path, char *const arg0[], ..., 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)**

NAME

fopen, fdopen – stream open functions

SYNOPSIS

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

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

fgetc, fgets, getc, getchar, gets, ungetc – input of characters and strings

SYNOPSIS

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

```
int fgetc(FILE *stream);
char *fgets(char *s, int size, FILE *stream);
int getc(FILE *stream);
int getchar(void);
char *gets(char *s);
int ungetc(int c, FILE *stream);
```

DESCRIPTION

fgetc() reads the next character from *stream* and returns it as an *unsigned char* cast to an *int*, or **EOF** on end of file or error.

getc() is equivalent to **fgetc()** except that it may be implemented as a macro which evaluates *stream* more than once.

getchar() is equivalent to **getc(stdin)**.

gets() reads a line from *stdin* into the buffer pointed to by *s* until either a terminating newline or **EOF**, which it replaces with **'\0'**. No check for buffer overrun is performed (see **BUGS** below).

fgets() reads in at most one less than *size* characters from *stream* and stores them into the buffer pointed to by *s*. Reading stops after an **EOF** or a newline. If a newline is read, it is stored into the buffer. A **'\0'** is stored after the last character in the buffer.

ungetc() pushes *c* back to *stream*, cast to *unsigned char*, where it is available for subsequent read operations. Pushed-back characters will be returned in reverse order; only one pushback is guaranteed.

Calls to the functions described here can be mixed with each other and with calls to other input functions from the *stdio* library for the same input stream.

For non-locking counterparts, see **unlocked_stdio(3)**.

RETURN VALUE

fgetc(), **getc()** and **getchar()** return the character read as an *unsigned char* cast to an *int* or **EOF** on end of file or error.

gets() and **fgets()** return *s* on success, and **NULL** on error or when end of file occurs while no characters have been read.

ungetc() returns *c* on success, or **EOF** on error.

CONFORMING TO

C89, C99. LSB deprecates **gets()**.

BUGS

Never use **gets()**. Because it is impossible to tell without knowing the data in advance how many characters **gets()** will read, and because **gets()** will continue to store characters past the end of the buffer, it is extremely dangerous to use. It has been used to break computer security. Use **fgets()** instead.

It is not advisable to mix calls to input functions from the *stdio* library with low-level calls to **read(2)** for the file descriptor associated with the input stream; the results will be undefined and very probably not what you want.

SEE ALSO

read(2), **write(2)**, **ferror(3)**, **fgetc(3)**, **fgetws(3)**, **fopen(3)**, **fread(3)**, **fseek(3)**, **getline(3)**, **getwchar(3)**, **puts(3)**, **scanf(3)**, **ungetc(3)**, **unlocked_stdio(3)**

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 an 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

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 *in6addr_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 *in6addr_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 sockaddr_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 *sockaddr_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 be changed to use *struct sockaddr_storage* for that instead.

SEE ALSO

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

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

```
#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. The currently understood formats are:

PF_INET ARPA Internet protocols

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

SOCK_STREAM
SOCK_DGRAM

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).

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**. 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.

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:

EACCES Permission to create a socket of the specified type and/or protocol is denied.
ENOMEM Insufficient user memory is available.

SEE ALSO

close(2), **read(2)**, **write(2)**, **accept(3N)**, **bind(3N)**, **connect(3N)**, **listen(3N)**,

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; /* blocksizes 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 */
};
```

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” blocksizes for efficient file system I/O. (Writing to a file in smaller chunks may cause an inefficient read-modify-rewrite.)

Not all of the Linux file systems implement all of the time fields. Some file system types allow mounting in such a way that file accesses do not cause an update of the *st_atime* field. (See "noatime" in `mount(8)`.)

The field *st_atime* is changed by file accesses, for example, by `execve(2)`, `mknod(2)`, `pipe(2)`, `utime(2)` and `read(2)` (of more than zero bytes). Other routines, like `mmap(2)`, may or may not update *st_atime*.

The field *st_mtime* is changed by file modifications, for example, by `mknod(2)`, `truncate(2)`, `utime(2)` and `write(2)` (of more than zero bytes). Moreover, *st_mtime* of a directory is changed by the creation or deletion of files in that directory. The *st_mtime* field is *not* changed for changes in owner, group, hard link count, or mode.

The field *st_ctime* is changed by writing or by setting inode information (i.e., owner, group, link count, mode, etc.).

The following POSIX macros are defined to check the file type using the *st_mode* field:

<code>S_ISREG(m)</code>	is it a regular file?
<code>S_ISDIR(m)</code>	directory?
<code>S_ISCHR(m)</code>	character device?
<code>S_ISBLK(m)</code>	block device?
<code>S_ISFIFO(m)</code>	FIFO (named pipe)?
<code>S_ISLNK(m)</code>	symbolic link? (Not in POSIX.1-1996.)
<code>S_ISSOCK(m)</code>	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)`

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	<code>waitpid()</code> 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	<code>waitpid()</code> 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)`