

accept(3)

accept(3)

**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 *ns*; 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.
- 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**

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

bind(3)

bind(3)

**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.
  - ENODIR** 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.

## NAME

fdopen – associate a stream with a file descriptor

## SYNOPSIS

```
#include <stdio.h>
FILE *fdopen(int fd, const char *mode);
```

## DESCRIPTION

The **fdopen()** function associates a stream with a file descriptor *fd*, whose value must be less than 255.

The *mode* argument is a character string having one of the following values:

<b>r</b> or <b>rb</b>	open a file for reading
<b>w</b> or <b>wb</b>	open a file for writing
<b>a</b> or <b>ab</b>	open a file for writing at end of file
<b>r+</b> or <b>rb+</b> or <b>r+b</b>	open a file for update (reading and writing)
<b>w+</b> or <b>wb+</b> or <b>w+b</b>	open a file for update (reading and writing)
<b>a+</b> or <b>ab+</b> or <b>a+b</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 *fd*.

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:

<b>EBADF</b>	The <i>fd</i> argument is not a valid file descriptor.
<b>EINVAL</b>	The <i>mode</i> argument is not a valid mode.
<b>EMFILE</b>	<b>FOPEN_MAX</b> streams are currently open in the calling process.
<b>EMFILE</b>	<b>STREAM_MAX</b> streams are currently open in the calling process.
<b>ENOMEM</b>	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

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**

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**

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:

<b>EACCESS</b>	Permission to create a socket of the specified type and/or protocol is denied.
<b>EMFILE</b>	The per-process descriptor table is full.
<b>ENOMEM</b>	Insufficient user memory is available.
<b>ENOSR</b>	There were insufficient STREAMS resources available to complete the operation.
<b>EPROTONOSUPPORT</b>	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)**

**NAME**

stat, fstat, lstat – get file status

**SYNOPSIS**

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

```
int stat(const char *fi_le_name, struct stat *buf);
int fstat(int fi_le_des, struct stat *buf);
int lstat(const char *fi_le_name, struct stat *buf);
```

**DESCRIPTION**

These functions return information about the specified file. You do not need any access rights to the file to get this information but you need search rights to all directories named in the path leading to the file.

**stat** stats the file pointed to by *fi\_le\_name* and fills in *buf*.

**lstat** is identical to **stat**, except in the case of a symbolic link, where the link itself is stat-ed, not the file that it refers to.

**fstat** is identical to **stat**, only the open file pointed to by *fi\_le\_des* (as returned by **open(2)**) is stat-ed in place of *fi\_le\_name*.

They all return a *stat* structure, which contains the following fields:

```
struct stat {
    dev_t    st_dev;    /* device */
    ino_t    st_ino;    /* inode */
    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 type (if inode device) */
    off_t    st_size;    /* total size, in bytes */
    blksize_t st_blksize; /* blocksize for filesystem 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 value *st\_size* gives the size of the file (if it is a regular file or a symlink) in bytes. The size of a symlink is the length of the pathname it contains, without trailing NUL.

The value *st\_blocks* gives the size of the file in 512-byte blocks. (This may be smaller than *st\_size*/512 e.g. when the file has holes.) The value *st\_blksize* gives the "preferred" blocksize 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 filesystems 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, e.g. 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, e.g. 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:

S_ISREG(m)	is it a regular file?
S_ISDIR(m)	directory?
S_ISCHR(m)	character device?
S_ISBLK(m)	block device?
S_ISFIFO(m)	fifo?
S_ISLNK(m)	symbolic link? (Not in POSIX.1-1996.)
S_ISSOCK(m)	socket? (Not in POSIX.1-1996.)

The following flags are defined for the *st\_mode* field:

S_IFMT	0170000	bitmask for the file type bitfields
S_IFSOCK	0140000	socket
S_IFLNK	0120000	symbolic link
S_IFREG	0100000	regular file
S_IFBLK	0060000	block device
S_IFDIR	0040000	directory
S_IFCHR	0020000	character device
S_IFIFO	0010000	fifo
S_ISUID	0004000	set UID bit
S_ISGID	0002000	set GID bit (see below)
S_ISVTX	0001000	sticky bit (see below)
S_IRWXU	00700	mask for file owner permissions
S_IRUSR	00400	owner has read permission
S_IWUSR	00200	owner has write permission
S_IXUSR	00100	owner has execute permission
S_IRWXG	00070	mask for group permissions
S_IRGRP	00040	group has read permission
S_IWGRP	00020	group has write permission
S_IXGRP	00010	group has execute permission
S_IRWXO	00007	mask for permissions for others (not in group)
S_IROTH	00004	others have read permission
S_IWOTH	00002	others have write permission
S_IXOTH	00001	others have execute permission

The set GID bit (S\_ISGID) has several special uses: For a directory it indicates that BSD semantics is to be used for that directory: files created there inherit their group ID from the directory, not from the effective group ID of the creating process, and directories created there will also get the S\_ISGID bit set. For a file that does not have the group execution bit (S\_IXGRP) set, it indicates mandatory file/record locking.

The 'sticky' bit (S\_ISVTX) on a directory means that a file in that directory can be renamed or deleted only by the owner of the file, by the owner of the directory, and by a privileged process.

**RETURN VALUE**

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

**SEE ALSO**

**chmod(2)**, **chown(2)**, **readlink(2)**, **utime(2)**, **capabilities(7)**

strerror(3)

strerror(3)

**NAME**

strerror, strerror\_r – return string describing error code

**SYNOPSIS**

#include <string.h>

char \*strerror(int *errnum*);  
int strerror\_r(int *errnum*, char \**buf*, size\_t *n*);

**DESCRIPTION**

The **strerror()** function returns a string describing the error code passed in the argument *errnum*, possibly using the LC\_MESSAGES part of the current locale to select the appropriate language. This string must not be modified by the application, but may be modified by a subsequent call to **perror()** or **strerror()**. No library function will modify this string.

The **strerror\_r()** function is similar to **strerror()**, but is thread safe. It returns the string in the user-supplied buffer *buf* of length *n*.

**RETURN VALUE**

The **strerror()** function returns the appropriate error description string, or an unknown error message if the error code is unknown. The value of *errno* is not changed for a successful call, and is set to a nonzero value upon error. The **strerror\_r()** function returns 0 on success and -1 on failure, setting *errno*.

**ERRORS**

**EINVAL**

The value of *errnum* is not a valid error number.

**ERANGE**

Insufficient storage was supplied to contain the error description string.

**CONFORMING TO**

SVID 3, POSIX, BSD 4.3, ISO/IEC 9899:1990 (C89).

**strerror\_r()** with prototype as given above is specified by SUSv3, and was in use under Digital Unix and HP Unix. An incompatible function, with prototype

char \*strerror\_r(int *errnum*, char \**buf*, size\_t *n*);

is a GNU extension used by glibc (since 2.0), and must be regarded as obsolete in view of SUSv3. The GNU version may, but need not, use the user-supplied buffer. If it does, the result may be truncated in case the supplied buffer is too small. The result is always NUL-terminated.

**SEE ALSO**

**errno(3)**, **perror(3)**, **strsignal(3)**

waitpid(2)

waitpid(2)

**NAME**

waitpid – wait for child process to change state

**SYNOPSIS**

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

pid\_t waitpid(pid\_t *pid*, int \**stat\_loc*, int *options*);

**DESCRIPTION**

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

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

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

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

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

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

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

**WCONTINUED**

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

**WNOHANG**

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

**WNOEXIT**

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)**