

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

fopen – open a stream

## SYNOPSIS

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

```
FILE *fopen(const char *filename, const char *mode);
```

## DESCRIPTION

The **fopen()** function opens the file whose pathname is the string pointed to by *filename*, and associates a stream with it.

The argument *mode* points to a string beginning with one of the following sequences:

<b>r</b> or <b>rb</b>	open file for reading
<b>w</b> or <b>wb</b>	truncate to zero length or create file for writing
<b>a</b> or <b>ab</b>	append; open or create file for writing at end-of-file
<b>r+</b> or <b>rb+</b> or <b>r+b</b>	open file for update (reading and writing)
<b>w+</b> or <b>wb+</b> or <b>w+b</b>	truncate to zero length or create file for update
<b>a+</b> or <b>ab+</b> or <b>a+b</b>	append; open or create file for update, writing at end-of-file

The character **b** has no effect, but is allowed for ISO C standard conformance. Opening a file with read mode (**r** as the first character in the *mode* argument) fails if the file does not exist or cannot be read.

When a file is opened with update mode (+ as the second or third character in the *mode* argument), both input and output may be performed on the associated stream. However, output must not be directly followed by input without an intervening call to **fflush(3S)** or to a file positioning function (**fseek(3S)**, **fsetpos(3S)** or **rewind(3S)**), and input must not be directly followed by output without an intervening call to a file positioning function, unless the input operation encounters end-of-file.

When opened, a stream is fully buffered if and only if it can be determined not to refer to an interactive device. The error and end-of-file indicators for the stream are cleared.

If *mode* is **w**, **a**, **w+** or **a+** and the file did not previously exist, upon successful completion, **fopen()** function will mark for update the **st\_atime**, **st\_ctime** and **st\_mtime** fields of the file and the **st\_ctime** and **st\_mtime** fields of the parent directory.

If *mode* is **w** or **w+** and the file did previously exist, upon successful completion, **fopen()** will mark for update the **st\_ctime** and **st\_mtime** fields of the file. The **fopen()** function will allocate a file descriptor as **open(2)** does.

The largest value that can be represented correctly in an object of type **off\_t** will be established as the offset maximum in the open file description.

## RETURN VALUES

Upon successful completion, **fopen()** returns a pointer to the object controlling the stream. Otherwise, a null pointer is returned, and **errno** is set to indicate the error.

**fopen()** may fail and not set **errno** if there are no free **stdio** streams.

## ERRORS

The **fopen()** function will fail if:

<b>EACCES</b>	Search permission is denied on a component of the path prefix, or the file exists and the permissions specified by <i>mode</i> are denied, or the file does not exist and write permission is denied for the parent directory of the file to be created.
<b>EINTR</b>	A signal was caught during <b>fopen()</b> .
<b>EISDIR</b>	The named file is a directory and <i>mode</i> requires write access.

## SEE ALSO

**fclose(3S)**, **fdopen(3S)**, **fflush(3S)**, **freopen(3S)**, **fsetpos(3S)**, **rewind(3S)**,

## NAME

fread, fwrite – binary stream input/output

## SYNOPSIS

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

```
size_t fread(void *ptr, size_t size, size_t nmemb, FILE *stream);
```

```
size_t fwrite(const void *ptr, size_t size, size_t nmemb, FILE *stream);
```

## DESCRIPTION

The function **fread** reads *nmemb* elements of data, each *size* bytes long, from the stream pointed to by *stream*, storing them at the location given by *ptr*.

The function **fwrite** writes *nmemb* elements of data, each *size* bytes long, to the stream pointed to by *stream*, obtaining them from the location given by *ptr*.

## RETURN VALUE

**fread** and **fwrite** return the number of items successfully read or written (i.e., not the number of characters).

## NAME

ip – Linux IPv4 protocol implementation

## SYNOPSIS

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

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

## DESCRIPTION

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

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

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

## ADDRESS FORMAT

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

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

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

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

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

## SEE ALSO

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

## NAME

calloc, malloc, free, realloc – Allocate and free dynamic memory

## SYNOPSIS

```
#include <stdlib.h>
```

```
void *calloc(size_t nmemb, size_t size);
void *malloc(size_t size);
void free(void *ptr);
void *realloc(void *ptr, size_t size);
```

## DESCRIPTION

`calloc()` allocates memory for an array of *nmemb* elements of *size* bytes each and returns a pointer to the allocated memory. The memory is set to zero.

`malloc()` allocates *size* bytes and returns a pointer to the allocated memory. The memory is not cleared.

`free()` frees the memory space pointed to by *ptr*, which must have been returned by a previous call to `malloc()`, `calloc()` or `realloc()`. Otherwise, or if `free(ptr)` has already been called before, undefined behaviour occurs. If *ptr* is `NULL`, no operation is performed.

`realloc()` changes the *size* of the memory block pointed to by *ptr* to *size* bytes. The contents will be unchanged to the minimum of the old and new sizes; newly allocated memory will be uninitialized. If *ptr* is `NULL`, the call is equivalent to `malloc(size)`; if *size* is equal to zero, the call is equivalent to `free(ptr)`. Unless *ptr* is `NULL`, it must have been returned by an earlier call to `malloc()`, `calloc()` or `realloc()`.

## RETURN VALUE

For `calloc()` and `malloc()`, the value returned is a pointer to the allocated memory, which is suitably aligned for any kind of variable, or `NULL` if the request fails.

`free()` returns no value.

`realloc()` returns a pointer to the newly allocated memory, which is suitably aligned for any kind of variable and may be different from *ptr*, or `NULL` if the request fails. If *size* was equal to 0, either `NULL` or a pointer suitable to be passed to `free()` is returned. If `realloc()` fails the original block is left untouched - it is not freed or moved.

## SEE ALSO

[brk\(2\)](#), [posix\\_memalign\(3\)](#)

**NAME**

sigaction – POSIX signal handling functions.

**SYNOPSIS**

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

```
int sigaction(int signum, const struct sigaction *act, struct sigaction *oldact);
```

**DESCRIPTION**

The **sigaction** system call is used to change the action taken by a process on receipt of a specific signal.

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

If *act* is non-null, the new action for signal *signum* is installed from *act*. If *oldact* is non-null, the previous action is saved in *oldact*.

The **sigaction** structure is defined as something like

```
struct sigaction {
    void (*sa_handler)(int);
    void (*sa_sigaction)(int, siginfo_t *, void *);
    sigset_t sa_mask;
    int sa_flags;
    void (*sa_restorer)(void);
}
```

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

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

*sa\_handler* specifies the action to be associated with *signum* and may be **SIG\_DFL** for the default action, **SIG\_IGN** to ignore this signal, or a pointer to a signal handling function.

*sa\_mask* gives a mask of signals which should be blocked during execution of the signal handler. In addition, the signal which triggered the handler will be blocked, unless the **SA\_NODEFER** or **SA\_NOMASK** flags are used.

*sa\_flags* specifies a set of flags which modify the behaviour of the signal handling process. It is formed by the bitwise OR of zero or more of the following:

**SA\_NOCLDSTOP**

If *signum* is **SIGCHLD**, do not receive notification when child processes stop (i.e., when child processes receive one of **SIGSTOP**, **SIGTSTP**, **SIGTTIN** or **SIGTTOU**).

**SA\_RESTART**

Provide behaviour compatible with BSD signal semantics by making certain system calls restartable across signals.

**RETURN VALUES**

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

**ERRORS****EINVAL**

An invalid signal was specified. This will also be generated if an attempt is made to change the action for **SIGKILL** or **SIGSTOP**, which cannot be caught.

**SEE ALSO**

**kill(1)**, **kill(2)**, **killpg(2)**, **pause(2)**, **sigsetops(3)**,

**NAME**

sigprocmask – change and/or examine caller's signal mask

**SYNOPSIS**

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

```
int sigprocmask(int how, const sigset_t *set, sigset_t *oset);
```

**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 to 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.

**SEE ALSO**

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

## NAME

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

## SYNOPSIS

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

## DESCRIPTION

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

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

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

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

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

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

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

## RETURN VALUES

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

## ERRORS

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

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

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

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

## SEE ALSO

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

## NAME

socket – create an endpoint for communication

## SYNOPSIS

```
cc [ flag ... ] file ... -lsocket -lnsl [ library ... ]
#include <sys/types.h>
#include <sys/socket.h>
int socket(int domain, int type, int protocol);
```

## DESCRIPTION

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

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

**PF\_UNIX** UNIX system internal protocols

**PF\_INET** ARPA Internet protocols

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

**SOCK\_STREAM**  
**SOCK\_DGRAM**  
**SOCK\_RAW**  
**SOCK\_SEQPACKET**  
**SOCK\_RDM**

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

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

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

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

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

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

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

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

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

#### RETURN VALUES

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

#### ERRORS

The **socket()** call fails if:

<b>EACCES</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 *file_name, struct stat *buf);
int lstat(const char *file_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 *file\_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.

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.)

#### RETURN VALUE

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

#### ERRORS

##### ENOENT

A component of the path *file\_name* does not exist, or the path is an empty string.

##### EACCES

Permission denied.

##### ENAMETOOLONG

File name too long.

#### SEE ALSO

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