

## NAME

connect – initiate a connection on a socket

## SYNOPSIS

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

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

## DESCRIPTION

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

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

## RETURN VALUE

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

## ERRORS

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

**EBADF**

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

**EFAULT**

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

**ENOTSOCK**

The file descriptor is not associated with a socket.

**EISCONN**

The socket is already connected.

**ECONNREFUSED**

No one listening on the remote address.

**ENETUNREACH**

Network is unreachable.

**EADDRINUSE**

Local address is already in use.

**EAFNOSUPPORT**

The passed address didn't have the correct address family in its *sa\_family* field.

**EACCES, EPERM**

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

## SEE ALSO

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

## NAME

opendir – open a directory / readdir – read a directory

## SYNOPSIS

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

```
#include <dirent.h>
```

```
DIR *opendir(const char *name);
```

```
struct dirent *readdir(DIR *dir);
```

```
int readdir_r(DIR *dirp, struct dirent *entry, struct dirent **result);
```

## DESCRIPTION opendir

The **opendir()** function opens a directory stream corresponding to the directory *name*, and returns a pointer to the directory stream. The stream is positioned at the first entry in the directory.

## RETURN VALUE

The **opendir()** function returns a pointer to the directory stream or NULL if an error occurred.

## DESCRIPTION readdir

The **readdir()** function returns a pointer to a dirent structure representing the next directory entry in the directory stream pointed to by *dir*. It returns NULL on reaching the end-of-file or if an error occurred.

## DESCRIPTION readdir\_r

The **readdir\_r()** function initializes the structure referenced by *entry* and stores a pointer to this structure in *result*. On successful return, the pointer returned at *result* will have the same value as the argument *entry*. Upon reaching the end of the directory stream, this pointer will have the value NULL.

The data returned by **readdir()** is overwritten by subsequent calls to **readdir()** for the **same** directory stream.

The *dirent* structure is defined as follows:

```
struct dirent {
    long    d_ino;           /* inode number */
    off_t   d_off;          /* offset to the next dirent */
    unsigned short d_reclen; /* length of this record */
    unsigned char d_type;    /* type of file */
    char     d_name[256];    /* filename */
};
```

## RETURN VALUE

The **readdir()** function returns a pointer to a dirent structure, or NULL if an error occurs or end-of-file is reached.

**readdir\_r()** returns 0 if successful or an error number to indicate failure.

## ERRORS

**EACCES**

Permission denied.

**ENOENT**

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

**ENOTDIR**

*name* is not a directory.

## 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

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 X3.159-1989 ("ANSI C").

## SEE ALSO

**open**(2), **stdio**(3), **unlocked\_stdio**(3)

**NAME**

gethostbyname – get network host entry

**SYNOPSIS**

```
#include <netdb.h>
extern int h_errno;
```

```
struct hostent *gethostbyname(const char *name);
```

**DESCRIPTION**

The `gethostbyname()` function returns a structure of type `hostent` for the given host `name`. Here `name` is either a host name, or an IPv4 address in standard dot notation, or an IPv6 address in colon (and possibly dot) notation. (See RFC 1884 for the description of IPv6 addresses.)

The `hostent` structure is defined in `<netdb.h>` as follows:

```
struct hostent {
    char    *h_name;        /* official name of host */
    char    **h_aliases;    /* alias list */
    int     h_addrtype;     /* host address type */
    int     h_length;       /* length of address */
    char    **h_addr_list;  /* list of addresses */
}
#define h_addr    h_addr_list[0] /* for backward compatibility */
```

The members of the `hostent` structure are:

`h_name`

The official name of the host.

`h_aliases`

A zero-terminated array of alternative names for the host.

`h_addrtype`

The type of address; always `AF_INET` at present.

`h_length`

The length of the address in bytes.

`h_addr_list`

A zero-terminated array of network addresses for the host in network byte order.

`h_addr`

The first address in `h_addr_list` for backward compatibility.

**RETURN VALUE**

The `gethostbyname()` function returns the `hostent` structure or a NULL pointer if an error occurs. On error, the `h_errno` variable holds an error number.

**ERRORS**

The variable `h_errno` can have the following values:

**HOST\_NOT\_FOUND**

The specified host is unknown.

**NO\_ADDRESS or NO\_DATA**

The requested name is valid but does not have an IP address.

**SEE ALSO**

`resolver(3)`, `hosts(5)`, `hostname(7)`, `resolv+(8)`, `named(8)`

**NAME**

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

**SYNOPSIS**

```
#include <stdio.h>
char *gets(char *s);
char *fgets(char *s, int n, FILE *stream);
int fputs(const char *s, FILE *stream);
int puts(const char *s);
```

**DESCRIPTION gets/fgets**

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

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

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

**RETURN VALUES**

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

**ERRORS**

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

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

**DESCRIPTION puts/fputs**

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

`puts()` writes the string `s` and a trailing newline to `stdout`.

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

**RETURN VALUE**

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

## 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

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.

## CONFORMING TO

ANSI-C

## SEE ALSO

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

**NAME**

memcpy – copy memory area

**SYNOPSIS**

```
#include <string.h>
```

```
void *memcpy(void *dest, const void *src, size_t n);
```

**DESCRIPTION**

The **memcpy()** function copies *n* bytes from memory area *src* to memory area *dest*. The memory areas may not overlap. Use **memmove(3)** if the memory areas do overlap.

**RETURN VALUE**

The **memcpy()** function returns a pointer to *dest*.

**CONFORMING TO**

SVID 3, BSD 4.3, ISO 9899

**SEE ALSO**

**bcopy(3)**, **memccpy(3)**, **memmove(3)**, **strcpy(3)**, **strncpy(3)**

**NAME**

pthread\_create – create a new thread / pthread\_exit – terminate the calling thread

**SYNOPSIS**

```
#include <pthread.h>
```

```
int pthread_create(pthread_t * thread, pthread_attr_t * attr, void * (*start_routine)(void *), void * arg);
```

```
void pthread_exit(void *retval);
```

**DESCRIPTION**

**pthread\_create** creates a new thread of control that executes concurrently with the calling thread. The new thread applies the function *start\_routine* passing it *arg* as first argument. The new thread terminates either explicitly, by calling **pthread\_exit(3)**, or implicitly, by returning from the *start\_routine* function. The latter case is equivalent to calling **pthread\_exit(3)** with the result returned by *start\_routine* as exit code.

The *attr* argument specifies thread attributes to be applied to the new thread. See **pthread\_attr\_init(3)** for a complete list of thread attributes. The *attr* argument can also be **NULL**, in which case default attributes are used: the created thread is joinable (not detached) and has default (non real-time) scheduling policy.

**pthread\_exit** terminates the execution of the calling thread. All cleanup handlers that have been set for the calling thread with **pthread\_cleanup\_push(3)** are executed in reverse order (the most recently pushed handler is executed first). Finalization functions for thread-specific data are then called for all keys that have non- **NULL** values associated with them in the calling thread (see **pthread\_key\_create(3)**). Finally, execution of the calling thread is stopped.

The *retval* argument is the return value of the thread. It can be consulted from another thread using **pthread\_join(3)**.

**RETURN VALUE**

On success, the identifier of the newly created thread is stored in the location pointed by the *thread* argument, and a 0 is returned. On error, a non-zero error code is returned.

The **pthread\_exit** function never returns.

**ERRORS****EAGAIN**

not enough system resources to create a process for the new thread.

**EAGAIN**

more than **PTHREAD\_THREADS\_MAX** threads are already active.

**AUTHOR**

Xavier Leroy <Xavier.Leroy@inria.fr>

**SEE ALSO**

**pthread\_join(3)**, **pthread\_detach(3)**, **pthread\_attr\_init(3)**.

**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 fstat(int fildes, 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.

**fstat** is identical to **stat**, only the open file pointed to by *fildes* (as returned by **open(2)**) is stat-ed in place of *file\_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)**

unlink(2)

unlink(2)

#### NAME

unlink – remove directory entry

#### SYNOPSIS

```
#include <unistd.h>
```

```
int unlink(const char *path);
```

#### DESCRIPTION

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

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

#### RETURN VALUES

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

#### ERRORS

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

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

#### SEE ALSO

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