

accept(2) accept(2)

NAME accept – accept a connection on a socket

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

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

DESCRIPTION

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

On success, these system calls return a nonnegative integer that is a file descriptor for the accepted socket. On error, `-1` is returned, and *errno* is set appropriately.

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)`

fdopen(3)

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:

- r** or **rb** open a file for reading
- w** or **wb** open a file for writing
- a** or **ab** open a file for writing at end of file
- r+** or **rb+** or **r+b** open a file for update (reading and writing)
- w+** or **wb+** or **w+b** open a file for update (reading and writing)
- a+** or **ab+** or **a+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:

- EBADF** The *fd* argument is not a valid file descriptor.
- EINVAL** The *mode* argument is not a valid mode.
- EMFILE** `FOPEN_MAX` streams are currently open in the calling process.
- EMFILE** `STREAM_MAX` streams are currently open in the calling process.
- ENOMEM** 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)`

flush(3) fflush(3)

NAME
flush — flush a stream

SYNOPSIS
#include <stdio.h>
int fflush(FILE *stream);

DESCRIPTION
For output streams, **fflush()** forces a write of all user-space buffered data for the given output or update *stream* via the stream's underlying write function.
For input streams associated with seekable files (e.g., disk files, but not pipes or terminals), **fflush()** discards any buffered data that has been fetched from the underlying file, but has not been consumed by the application.
The open status of the stream is unaffected.
If the *stream* argument is NULL, **fflush()** flushes *all* open output streams.
For a nonlocking counterpart, see **unlocked_stdio(3)**.

RETURN VALUE
Upon successful completion 0 is returned. Otherwise, **EOF** is returned and *errno* is set to indicate the error.

ERRORS
EBADF
stream is not an open stream, or is not open for writing.
The function **fflush()** may also fail and set *errno* for any of the errors specified for **write(2)**.

SEE ALSO
fsync(2), **sync(2)**, **write(2)**, **fclose(3)**, **fileno(3)**, **fopen(3)**, **setbuf(3)**, **unlocked_stdio(3)**

FNMATCH(3) FNMATCH(3)

NAME
fnmatch — match filename or pathname

SYNOPSIS
#include <fnmatch.h>
int fnmatch(const char *pattern, const char *string, int flags);

DESCRIPTION
The **fnmatch()** function checks whether the *string* argument matches the *pattern* argument, which is a shell wildcard pattern.
The *flags* argument modifies the behavior; it is the bitwise OR of zero or more of the following flags:
FNM_NOESCAPE
If this flag is set, treat backslash as an ordinary character, instead of an escape character.
FNM_PATHNAME
If this flag is set, match a slash in *string* only with a slash in *pattern* and not by an asterisk (*) or a question mark (?) metacharacter, not by a bracket expression ([]) containing a slash.
FNM_PERIOD
If this flag is set, a leading period in *string* has to be matched exactly by a period in *pattern*. A period is considered to be leading if it is the first character in *string*, or if both **FNM_PATHNAME** and **FNM_PERIOD** is set and the period immediately follows a slash.
FNM_FILE_NAME
This is a GNU synonym for **FNM_PATHNAME**.
FNM_LEADING_DIR
If this flag (a GNU extension) is set, the pattern is considered to be matched if it matches an initial segment of *string* which is followed by a slash. This flag is mainly for the internal use of *glibc* and is only implemented in certain cases.
FNM_CASEFOLD
If this flag (a GNU extension) is set, the pattern is matched case-insensitively.

RETURN VALUE
Zero if *string* matches *pattern*, **FNM_NOMATCH** if there is no match or another nonzero value if there is an error.

CONFORMING TO
POSIX.2. The **FNM_FILE_NAME**, **FNM_LEADING_DIR**, and **FNM_CASEFOLD** flags are GNU extensions.

flush(3) fflush(3)

NAME
flush — flush a stream

SYNOPSIS
#include <stdio.h>
int fflush(FILE *stream);

DESCRIPTION
For output streams, **fflush()** forces a write of all user-space buffered data for the given output or update *stream* via the stream's underlying write function.
For input streams associated with seekable files (e.g., disk files, but not pipes or terminals), **fflush()** discards any buffered data that has been fetched from the underlying file, but has not been consumed by the application.
The open status of the stream is unaffected.
If the *stream* argument is NULL, **fflush()** flushes *all* open output streams.
For a nonlocking counterpart, see **unlocked_stdio(3)**.

RETURN VALUE
Upon successful completion 0 is returned. Otherwise, **EOF** is returned and *errno* is set to indicate the error.

ERRORS
EBADF
stream is not an open stream, or is not open for writing.
The function **fflush()** may also fail and set *errno* for any of the errors specified for **write(2)**.

SEE ALSO
fsync(2), **sync(2)**, **write(2)**, **fclose(3)**, **fileno(3)**, **fopen(3)**, **setbuf(3)**, **unlocked_stdio(3)**

opendir/readdir(3) opendir/readdir(3)

NAME
 opendir – open a directory / readdir – read a directory

SYNOPSIS

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

DIR *opendir(const char *name);
int closedir(DIR *dirp);
struct dirent *readdir(DIR *dir);
```

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. On error, NULL is returned, and *errno* is set appropriately.

DESCRIPTION *closedir*
 The **closedir()** function closes the directory stream associated with *dirp*. A successful call to **closedir()** also closes the underlying file descriptor associated with *dirp*. The directory stream descriptor *dirp* is not available after this call.

RETURN VALUE
 The **closedir()** function returns 0 on success. On error, -1 is returned, and *errno* is set appropriately.

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. It is safe to use **readdir()** inside threads if the pointers passed as *dir* are created by distinct calls to **opendir()**. 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 */
    char    d_name[256]; /* filename */
};
```

RETURN VALUE
 On success, **readdir()** returns a pointer to a *dirent* structure. (This structure may be statically allocated; do not attempt to **free(3)** it.)

If the end of the directory stream is reached, NULL is returned and *errno* is not changed. If an error occurs, NULL is returned and *errno* is set appropriately. To distinguish end of stream and from an error, set *errno* to zero before calling **readdir()** and then check the value of *errno* if NULL is returned.

ERRORS
EACCES Permission denied.
ENOENT Directory does not exist, or *name* is an empty string.
ENOTDIR *name* is not a directory.

ipv6/socket(7) ipv6/socket(7)

NAME
 ipv6, AF_INET6 – Linux IPv6 protocol implementation

SYNOPSIS

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

tcp6_socket = socket(AF_INET6, SOCK_STREAM, 0);
raw6_socket = socket(AF_INET6, SOCK_RAW, protocol);
udp6_socket = socket(AF_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.

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 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)**)

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

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)**

stat(2) stat(2) stat(2) stat(2) stat(2)

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; /* blocksize 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" 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 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 "noinst" 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:

- 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 (named pipe)?
- S_ISLNK(m)** symbolic link? (Not in POSIX.1-1996.)
- S_ISSOCK(m)** 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)**

strtok(3) strtok(3)

NAME
 strtok, strtok_r – extract tokens from strings

SYNOPSIS

```
#include <string.h>

char *strtok(char *str, const char *delim);
char *strtok_r(char *str, const char *delim, char **saveptr);
```

DESCRIPTION
 The **strtok()** function breaks a string into a sequence of zero or more nonempty tokens. On the first call to **strtok()** the string to be parsed should be specified in *str*. In each subsequent call that should parse the same string, *str* must be NULL.

The *delim* argument specifies a set of bytes that delimit the tokens in the parsed string. The caller may specify different strings in *delim* in successive calls that parse the same string.

Each call to **strtok()** returns a pointer to a null-terminated string containing the next token. This string does not include the delimiting byte. If no more tokens are found, **strtok()** returns NULL.

A sequence of calls to **strtok()** that operate on the same string maintains a pointer that determines the point from which to start searching for the next token. The first call to **strtok()** sets this pointer to point to the first byte of the string. The start of the next token is determined by scanning forward for the next nondelimiter byte in *str*. If such a byte is found, it is taken as the start of the next token. If no such byte is found, then there are no more tokens, and **strtok()** returns NULL. (A string that is empty or that contains only delimiters will thus cause **strtok()** to return NULL on the first call.)

The end of each token is found by scanning forward until either the next delimiter byte is found or until the terminating null byte (‘\0’) is encountered. If a delimiter byte is found, it is overwritten with a null byte to terminate the current token, and **strtok()** saves a pointer to the following byte; that pointer will be used as the starting point when searching for the next token. In this case, **strtok()** returns a pointer to the start of the found token.

From the above description, it follows that a sequence of two or more contiguous delimiter bytes in the parsed string is considered to be a single delimiter, and that delimiter bytes at the start or end of the string are ignored. Put another way: the tokens returned by **strtok()** are always nonempty strings. Thus, for example, given the string "aaa;bbb", successive calls to **strtok()** that specify the delimiter string ";" would return the strings "aaa" and "bbb", and then a null pointer.

The **strtok_r()** function is a reentrant version **strtok()**. The *saveptr* argument is a pointer to a *char ** variable that is used internally by **strtok_r()** in order to maintain context between successive calls that parse the same string. On the first call to **strtok_r()**, *str* should point to the string to be parsed, and the value of *saveptr* is ignored. In subsequent calls, *str* should be NULL, and *saveptr* should be unchanged since the previous call.

Different strings may be parsed concurrently using sequences of calls to **strtok_r()** that specify different *saveptr* arguments.

RETURN VALUE
strtok() and **strtok_r()** return a pointer to the next token, or NULL if there are no more tokens.

ATTRIBUTES
Multithreading (see pthreads(7))
 The **strtok()** function is not thread-safe, the **strtok_r()** function is thread-safe.

strerror(3) strerror(3)

NAME
 strerror — get error message string

SYNOPSIS

```
#include <string.h>

const char *strerror(int errnum);
```

DESCRIPTION
 The functionality described on this reference page is aligned with the ISO C standard. Any conflict between the requirements described here and the ISO C standard is unintentional. This volume of POSIX.1-2017 defers to the ISO C standard.

The *strerror()* function shall map the error number in *errnum* to a locale-dependent error message string and shall return a pointer to it. Typically, the values for *errnum* come from *errno*, but *strerror()* shall map any value of type **int** to a message.

The application shall not modify the string returned. The returned string pointer might be invalidated or the string content might be overwritten by a subsequent call to *strerror()* in the same thread. The returned pointer and the string content might also be invalidated if the calling thread is terminated.

The contents of the error message strings returned by *strerror()* should be determined by the setting of the *LC_MESSAGES* category in the current locale.

The implementation shall behave as if no function defined in this volume of POSIX.1-2017 calls *strerror()*.

The *strerror()* function shall not change the setting of *errno* if successful.

If the value of *errnum* is a valid error number, the message string shall indicate what error occurred; if the value of *errnum* is zero, the message string shall either be an empty string or indicate that no error occurred; otherwise, if these functions complete successfully, the message string shall indicate that an unknown error occurred.

RETURN VALUE
 Upon completion, whether successful or not, *strerror()* shall return a pointer to the generated message string. On error *errno* may be set, but no return value is reserved to indicate an error.

Upon successful completion, *strerror_0()* shall return a pointer to the generated message string. If *errnum* is not a valid error number, *errno* may be set to [EINVAL], but a pointer to a message string shall still be returned. If any other error occurs, *errno* shall be set to indicate the error and a null pointer shall be returned.

Upon successful completion, *strerror_r()* shall return 0. Otherwise, an error number shall be returned to indicate the error.

ERRORS
 This function may fail if:

EINVAL
 The value of *errnum* is neither a valid error number nor zero. *The following sections are informative.*

EXAMPLES
 None.

SEE ALSO
pthread_0()

The Base Definitions volume of POSIX.1-2017, **<string.h>**