

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 *s*; 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 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.

**EWOLDBLOCK** 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(2)

bind(2)

**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.

**ENOSOCK** *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.

**ENOTDIR** A component of the path prefix of the pathname in *name* is not a directory.

**EROFS** The inode would reside on a read-only file system.

#### 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.

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);
```

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

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

**DESCRIPTION**

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.

feof/ferror/fileno(3)

feof/ferror/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 C89 and C99.

**SEE ALSO**

`open(2)`, `fdopen(3)`, `stdio(3)`, `unlocked_stdio(3)`

fopen/fdopen/filen0(3)

fopen/fdopen/filen0(3)

#### NAME

fopen, fdopen, fileno – stream open functions

#### SYNOPSIS

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

```
FILE *fopen(const char *path, const char *mode);  
FILE *fdopen(int fd, const char *mode);  
int fileno(FILE *stream);
```

#### 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, *fd*. 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 *fd*, 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.

The function **fileno()** examines the argument *stream* and returns its integer descriptor.

#### 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 **open** 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)**

fread/fwrite(3)

fread/fwrite(3)

#### 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*.

For nonlocking counterparts, see **unlocked\_stdio(3)**.

#### RETURN VALUE

**fread()** and **fwrite()** return the number of items successfully read or written (i.e., not the number of characters). If an error occurs, or the end-of-file is reached, the return value is a short item count (or zero).

**fread()** does not distinguish between end-of-file and error, and callers must use **feof(3)** and **ferror(3)** to determine which occurred.

#### CONFORMING TO

C89, POSIX.1-2001.

#### SEE ALSO

**read(2)**, **write(2)**, **feof(3)**, **ferror(3)**, **unlocked\_stdio(3)**

socket(2) / ipv6(7)

socket(2) / ipv6(7)

listen(2)

listen(2)

**NAME**

ipv6, PF\_INET6 – Linux IPv6 protocol implementation

**SYNOPSIS**

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

```
tcp6_socket = socket(PF_INET6, SOCK_STREAM, 0);
raw6_socket = socket(PF_INET6, SOCK_RAW, protocol);
udp6_socket = socket(PF_INET6, SOCK_DGRAM, protocol);
```

**DESCRIPTION**

Linux 2.2 optionally implements the Internet Protocol version 6. This man page contains a description of the IPv6 basic API as implemented by the Linux kernel and glibc 2.1. The interface is based on the BSD sockets interface; see [socket\(7\)](#).

The IPv6 API aims to be mostly compatible with the [ip\(7\)](#) v4 API. Only differences are described in this man page.

To bind an [AF\\_INET6](#) socket to any process the local address should be copied from the *inaddr\_any* variable which has *in6\_addr* type. In static initializations [IN6ADDR\\_ANY\\_INIT](#) may also be used, which expands to a constant expression. Both of them are in network order.

The IPv6 loopback address (:::1) is available in the global *inaddr\_loopback* variable. For initializations [IN6ADDR\\_LOOPBACK\\_INIT](#) should be used.

IPv4 connections can be handled with the v6 API by using the v4-mapped-on-v6 address type; thus a program only needs only to support this API type to support both protocols. This is handled transparently by the address handling functions in libc.

IPv4 and IPv6 share the local port space. When you get an IPv4 connection or packet to a IPv6 socket its source address will be mapped to v6 and it will be mapped to v6.

**Address Format**

```
struct sockadd_in6 {
    uint6_t    sin6_family; /* AF_INET6 */
    uint6_t    sin6_port; /* port number */
    uint32_t   sin6_flowinfo; /* IPv6 flow information */
    struct in6_addr sin6_addr; /* IPv6 address */
    uint32_t   sin6_scope_id; /* Scope ID (new in 2.4) */
};

struct in6_addr {
    unsigned char s6_addr[16]; /* IPv6 address */
};
```

*sin6\_family* is always set to [AF\\_INET6](#); *sin6\_port* is the protocol port (see [sin\\_port](#) in [ip\(7\)](#)); *sin6\_flowinfo* is the IPv6 flow identifier; *sin6\_addr* is the 128-bit IPv6 address. *sin6\_scope\_id* is an ID of depending of on the scope of the address. It is new in Linux 2.4. Linux only supports it for link scope addresses, in that case *sin6\_scope\_id* contains the interface index (see [netdevice\(7\)](#))

**NOTES**

The *sockadd\_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 use *struct sockaddr\_storage* for that instead.

**SEE ALSO**

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

**NAME**

listen – listen for connections on a socket

**SYNOPSIS**

```
#include <sys/types.h> /* See NOTES */
#include <sys/socket.h>
```

```
int listen(int sockfd, int backlog);
```

**DESCRIPTION**

[listen\(\)](#) marks the socket referred to by *sockfd* as a passive socket, that is, as a socket that will be used to accept incoming connection requests using [accept\(2\)](#).

The *sockfd* argument is a file descriptor that refers to a socket of type [SOCK\\_STREAM](#) or [SOCK\\_SEQPACKET](#).

The *backlog* argument defines the maximum length to which the queue of pending connections for *sockfd* may grow. If a connection request arrives when the queue is full, the client may receive an error with an indication of [ECONNREFUSED](#) or, if the underlying protocol supports retransmission, the request may be ignored so that a later reattempt at connection succeeds.

**RETURN VALUE**

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

**ERRORS**

**EADDRINUSE**

Another socket is already listening on the same port.

**EBADF**

The argument *sockfd* is not a valid descriptor.

**ENOTSOCK**

The argument *sockfd* is not a socket.

**NOTES**

To accept connections, the following steps are performed:

1. A socket is created with [socket\(2\)](#).
2. The socket is bound to a local address using [bind\(2\)](#), so that other sockets may be [connect\(2\)](#)ed to it.
3. A willingness to accept incoming connections and a queue limit for incoming connections are specified with [listen\(\)](#).
4. Connections are accepted with [accept\(2\)](#).

If the *backlog* argument is greater than the value in */proc/sys/net/core/somaxconn*, then it is silently truncated to that value; the default value in this file is 128.

**EXAMPLE**

See [bind\(2\)](#).

**SEE ALSO**

[accept\(2\)](#), [bind\(2\)](#), [connect\(2\)](#), [socket\(2\)](#), [socket\(7\)](#)

pthread\_create(pthread\_t(3))

pthread\_create(pthread\_exit(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 <XavierLeroy@inria.fr>

#### SEE ALSO

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

pthread\_join(3)

pthread\_join(3)

#### NAME

pthread\_join – join with a terminated thread

#### SYNOPSIS

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

```
int pthread_join(pthread_t t, void **retval);
```

```
Compile and link with -pthread.
```

#### DESCRIPTION

The **pthread\_join(3)** function waits for the thread specified by *thread* to terminate. If that thread has already terminated, then **pthread\_join(3)** returns immediately. The thread specified by *thread* must be joinable.

If *retval* is not **NULL**, then **pthread\_join(3)** copies the exit status of the target thread (i.e., the value that the target thread supplied to **pthread\_exit(3)**) into the location pointed to by *retval*. If the target thread was canceled, then **PTHREAD\_CANCELED** is placed in *retval*.

If multiple threads simultaneously try to join with the same thread, the results are undefined. If the thread calling **pthread\_join(3)** is canceled, then the target thread will remain joinable (i.e., it will not be detached).

#### RETURN VALUE

On success, **pthread\_join(3)** returns 0; on error, it returns an error number.

#### ERRORS

##### EDEADLK

A deadlock was detected (e.g., two threads tried to join with each other), or *thread* specifies the calling thread.

##### EINVAL

*thread* is not a joinable thread.

##### EINVAL

Another thread is already waiting to join with this thread.

##### ESRCH

No thread with the ID *thread* could be found.

#### NOTES

After a successful call to **pthread\_join(3)**, the caller is guaranteed that the target thread has terminated.

Joining with a thread that has previously been joined results in undefined behavior.

Failure to join with a thread that is joinable (i.e., one that is not detached), produces a "zombie thread". Avoid doing this, since each zombie thread consumes some system resources, and when enough zombie threads have accumulated, it will no longer be possible to create new threads (or processes).

There is no pthreads analog of *waitpid(1)*, *killatime(0)*, that is, "join with any terminated thread". If you believe you need this functionality, you probably need to rethink your application design.

All of the threads in a process are peers: any thread can join with any other thread in the process.

#### EXAMPLE

See **pthread\_create(3)**.

#### SEE ALSO

**pthread\_cancel(3)**, **pthread\_create(3)**, **pthread\_detach(3)**, **pthread\_exit(3)**, **pthread\_tryjoin\_np(3)**, **pthread\_t(7)**