

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 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(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.
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.
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.

dup(2)

dup(2)

#### NAME

dup, dup2 – duplicate a file descriptor

#### SYNOPSIS

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

```
int dup(int oldfd);
int dup2(int oldfd, int newfd);
```

#### DESCRIPTION

**dup()** and **dup2()** create a copy of the file descriptor *oldfd*.

**dup()** uses the lowest-numbered unused descriptor for the new descriptor.

**dup2()** makes *newfd* be the copy of *oldfd*, closing *newfd* first if necessary, but note the following:

- \* If *oldfd* is not a valid file descriptor, then the call fails, and *newfd* is not closed.
- \* If *oldfd* is a valid file descriptor, and *newfd* has the same value as *oldfd*, then **dup2()** does nothing, and returns *newfd*.

After a successful return from **dup()** or **dup2()**, the old and new file descriptors may be used interchangeably. They refer to the same open file description (see **open(2)**) and thus share file offset and file status flags; for example, if the file offset is modified by using **lseek(2)** on one of the descriptors, the offset is also changed for the other.

The two descriptors do not share file descriptor flags (the close-on-exec flag). The close-on-exec flag (**FD\_CLOEXEC**; see **fcntl(2)**) for the duplicate descriptor is off.

#### RETURN VALUE

**dup()** and **dup2()** return the new descriptor, or  $-1$  if an error occurred (in which case, *errno* is set appropriately).

#### ERRORS

##### EBADF

*oldfd* isn't an open file descriptor, or *newfd* is out of the allowed range for file descriptors.

##### EBUSY

(Linux only) This may be returned by **dup2()** during a race condition with **open(2)** and **dup()**.

##### EINTR

The **dup2()** call was interrupted by a signal; see **signal(7)**.

##### EMFILE

The process already has the maximum number of file descriptors open and tried to open a new one.

#### NOTES

The error returned by **dup2()** is different from that returned by **fcntl(..., F\_DUPFD, ...)** when *newfd* is out of range. On some systems **dup2()** also sometimes returns **EINVAL** like **F\_DUPFD**.

If *newfd* was open, any errors that would have been reported at **close(2)** time are lost. A careful programmer will not use **dup2()** without closing *newfd* first.

#### SEE ALSO

**close(2)**, **fcntl(2)**, **open(2)**

fopen/fdopen/fileno(3)

fopen/fdopen/fileno(3)

#### NAME

fopen, fdopen, fileno – stream open functions

#### SYNOPSIS

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

```
FILE *fopen(const char *path, const char *mode);
FILE *fdopen(int fildes, 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, *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.

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 **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)**

getc/fgets/putc/fputs(3)

getc/fgets/putc/fputs(3)

#### NAME

fgetc, fgets, getc, getchar, fputc, fputs, putc, putchar – input and output of characters and strings

#### SYNOPSIS

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

```
int fgetc(FILE *stream);
char *fgets(char *s, int size, FILE *stream);
int getc(FILE *stream);
int getchar(void);
int fputc(int c, FILE *stream);
int fputs(const char *s, FILE *stream);
int putc(int c, FILE *stream);
int putchar(int c);
```

#### DESCRIPTION

**fgetc()** reads the next character from *stream* and returns it as an *unsigned char* cast to an *int*, or **EOF** on end of file or error.

**getc()** is equivalent to **fgetc()** except that it may be implemented as a macro which evaluates *stream* more than once.

**getchar()** is equivalent to **getc(stdin)**.

**fgets()** reads in at most one less than *size* characters from *stream* and stores them into the buffer pointed to by *s*. Reading stops after an **EOF** or a newline. If a newline is read, it is stored into the buffer. A **'0'** is stored after the last character in the buffer.

**fputc()** writes the character *c*, cast to an *unsigned char*, to *stream*.

**fputs()** writes the string *s* to *stream*, without its terminating null byte (**'0'**).

**putc()** is equivalent to **fputc()** except that it may be implemented as a macro which evaluates *stream* more than once.

**putchar(c)**; is equivalent to **putc(c, 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

**fgetc()**, **getc()** and **getchar()** return the character read as an *unsigned char* cast to an *int* or **EOF** on end of file or error.

**fgets()** returns *s* on success, and **NULL** on error or when end of file occurs while no characters have been read. **fputc()**, **putc()** and **putchar()** return the character written as an *unsigned char* cast to an *int* or **EOF** on error.

**fputs()** returns a nonnegative number on success, or **EOF** on error.

#### SEE ALSO

**read(2)**, **write(2)**, **ferror(3)**, **fgetc(3)**, **fgets(3)**, **fopen(3)**, **fread(3)**, **fseek(3)**, **getline(3)**, **getwchar(3)**, **scanf(3)**, **ungetc(3)**, **write(2)**, **ferror(3)**, **fopen(3)**, **fputc(3)**, **fputws(3)**, **fseek(3)**, **fwrite(3)**, **gets(3)**, **putwchar(3)**, **scanf(3)**, **unlocked\_stdio(3)**

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)**

listen(2)

listen(2)

#### 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)**

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);
```

#### 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. 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 */
    off_t       d_off;        /* offset to the next dirent */
    unsigned short d_reclen;  /* length of this record */
    unsigned char d_type;     /* type of file; not supported by all filesystem types */
    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.

#### ERRORS

##### EACCES

Permission denied.

##### ENOENT

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

##### ENOTDIR

*name* is not a directory.

pthread\_create/pthread\_exit(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 <Xavier.Leroy@inria.fr>

#### SEE ALSO

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

strdup(3)

strdup(3)

#### NAME

strdup, strdup – duplicate a string

#### SYNOPSIS

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

```
char *strdup(const char *s);
```

```
char *strndup(const char *s, size_t n);
```

#### DESCRIPTION

The **strdup**() function returns a pointer to a new string which is a duplicate of the string *s*. Memory for the new string is obtained with **malloc**(3), and can be freed with **free**(3).

The **strndup**() function is similar, but copies at most *n* bytes. If *s* is longer than *n*, only *n* bytes are copied, and a terminating null byte ('\0') is added.

#### RETURN VALUE

On success, the **strdup**() function returns a pointer to the duplicated string. It returns **NULL** if insufficient memory was available, with *errno* set to indicate the cause of the error.

#### ERRORS

##### ENOMEM

Insufficient memory available to allocate duplicate string.

#### CONFORMING TO

**strdup**() conforms to SVr4, 4.3BSD, POSIX.1-2001. **strndup**() conforms to POSIX.1-2008.

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.