

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 <i>s</i> 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)**

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	<i>s</i> is not a valid descriptor.
EINVAL	<i>namelen</i> 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	<i>s</i> 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 <i>name</i> .
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 <i>name</i> .
ENOENT	A component of the path prefix of the pathname in <i>name</i> does not exist.
ENODIR	A component of the path prefix of the pathname in <i>name</i> is not a directory.
EROFS	The inode would reside on a read-only file system.

SEE ALSO

unlink(2), **socket(3N)**, **attributes(5)**, **socket(5)**

NOTES

Binding a name in the UNIX domain creates a socket in the file system that must be deleted by the caller when it is no longer needed (using **unlink(2)**).

The rules used in name binding vary between communication domains.

NAME

fdopen – associate a stream with a file descriptor

SYNOPSIS

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

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

DESCRIPTION

The **fdopen()** function associates a stream with a file descriptor *fildes*, 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 *fildes*.

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 <i>fildes</i> argument is not a valid file descriptor.
EINVAL	The <i>mode</i> 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)**

NAME

ip – Linux IPv4 protocol implementation

SYNOPSIS

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

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

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

DESCRIPTION

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

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

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

ADDRESS FORMAT

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

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

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

sin_addr is the IP host address. The *addr* member of **struct in_addr** contains the host interface address in network order. **in_addr** should be only accessed using the **inet_aton(3)**, **inet_addr(3)**, **inet_makeaddr(3)** library functions or directly with the name resolver (see **gethostbyname(3)**).

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

SEE ALSO

sendmsg(2), **recvmsg(2)**, **socket(7)**, **netlink(7)**, **tcp(7)**, **udp(7)**, **raw(7)**, **ipfw(7)**

NAME

opendir – open a directory

SYNOPSIS

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

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

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

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.

ERRORS**EACCES**

Permission denied.

EMFILE

Too many file descriptors in use by process.

ENFILE

Too many files are currently open in the system.

ENOENT

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

ENOMEM

Insufficient memory to complete the operation.

ENOTDIR

name is not a directory.

CONFORMING TO

SVID 3, POSIX, BSD 4.3

SEE ALSO

open(2), **readdir(3)**, **closedir(3)**, **rewinddir(3)**, **seekdir(3)**, **telldir(3)**, **scandir(3)**

NAME

readdir – read a directory

SYNOPSIS

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

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

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

DESCRIPTION

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.

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.

ERRORS**EBADF**

Invalid directory stream descriptor *dir*.

CONFORMING TO

SVID 3, POSIX, BSD 4.3

According to POSIX, the *dirent* structure contains a field *char d_name[]* of unspecified size, with at most **NAME_MAX** characters preceding the terminating null character. Use of other fields will harm the portability of your programs.

BUGS

Field *d_type* is not implemented as of libc6 2.1 and will always return DT_UNKNOWN (0).

SEE ALSO

read(2), **opendir(3)**, **closedir(3)**, **rewinddir(3)**, **seekdir(3)**, **telldir(3)**, **scandir(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:

EACCES	Permission to create a socket of the specified type and/or protocol is denied.
EMFILE	The per-process descriptor table is full.
ENOMEM	Insufficient user memory is available.
ENOSR	There were insufficient STREAMS resources available to complete the operation.
EPROTONOSUPPORT	The protocol type or the specified protocol is not supported within this domain.

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

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)**