

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

fork – create a child process

SYNOPSIS

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

```
pid_t fork(void);
```

DESCRIPTION

fork() creates a child process that differs from the parent process only in its PID and PPID, and in the fact that resource utilizations are set to 0. File locks and pending signals are not inherited.

Under Linux, **fork()** is implemented using copy-on-write pages, so the only penalty that it incurs is the time and memory required to duplicate the parent's page tables, and to create a unique task structure for the child.

RETURN VALUE

On success, the PID of the child process is returned in the parent's thread of execution, and a 0 is returned in the child's thread of execution. On failure, a **-1** will be returned in the parent's context, no child process will be created, and *errno* will be set appropriately.

ERRORS

EAGAIN

fork() cannot allocate sufficient memory to copy the parent's page tables and allocate a task structure for the child.

EAGAIN

It was not possible to create a new process because the caller's **RLIMIT_NPROC** resource limit was encountered. To exceed this limit, the process must have either the **CAP_SYS_ADMIN** or the **CAP_SYS_RESOURCE** capability.

ENOMEM

fork() failed to allocate the necessary kernel structures because memory is tight.

CONFORMING TO

The **fork()** call conforms to SVr4, SVID, POSIX, X/OPEN, 4.3BSD.

EXAMPLE

See **pipe(2)** and **wait(2)**.

SEE ALSO

clone(2), **execve(2)**, **setrlimit(2)**, **unshare(2)**, **vfork(2)**, **wait(2)**, **capabilities(7)**

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

sleep – Sleep for the specified number of seconds

SYNOPSIS

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

```
unsigned int sleep(unsigned int seconds);
```

DESCRIPTION

[sleep\(\)](#) makes the current process sleep until *seconds* seconds have elapsed or a signal arrives which is not ignored.

RETURN VALUE

Zero if the requested time has elapsed, or the number of seconds left to sleep.

CONFORMING TO

POSIX.1

BUGS

[sleep\(\)](#) may be implemented using **SIGALRM**; mixing calls to [alarm\(\)](#) and [sleep\(\)](#) is a bad idea.

Using [longjmp\(\)](#) from a signal handler or modifying the handling of **SIGALRM** while sleeping will cause undefined results.

SEE ALSO

[alarm\(2\)](#), [signal\(2\)](#)