

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)**

bbuffer(3)

bbuffer(3)

**NAME**

bbCreate, bbPut, bbGet, bbDestroy – A synchronized bounded-buffer implementation

**SYNOPSIS**

```
#include "bbuffer.h"
```

```
BNDBUF *bbCreate(size_t size);
void bbPut(BNDBUF * bb, void * value);
void* bbGet(BNDBUF * bb);
void bbDestroy(BNDBUF * bb);
```

**DESCRIPTION**

Bounded-buffer implementation of a FIFO queue. Manages **void\*** and supports multiple concurrent readers and writers. Provides the following functions:

**bbCreate()** creates a new bounded buffer for up to *size* elements. If an error occurs during the initialization, the implementation frees all resources already allocated by then and returns **NULL**.

**bbPut()** stores the *value* in the bounded buffer. If the buffer is full (i.e., it currently contains *size* elements), the call to **bbPut()** blocks until the value can be stored.

**bbGet()** returns the next value from the bounded buffer. If the buffer is empty, the call blocks until a value is available.

Both **bbPut()** and **bbGet()** are synchronized internally and thus can be called concurrently without the need for further synchronization.

**bbDestroy()** releases any resources related to the bounded buffer itself. It does not call **free()** on the elements stored in the buffer.

**RETURN VALUE**

**bbCreate()** returns a pointer to the allocated bounded buffer, or **NULL** if the request fails.

**bbPut()** returns no value.

**bbGet()** returns the next value stored in the bounded buffer.

**bbDestroy()** returns no value.

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)**

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/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)**

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.

#### SEE ALSO

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

## NAME

scanf, fscanf, sscanf – input format conversion

## SYNOPSIS

```
#include <stdio.h>

int scanf(const char *format, ...);
int fscanf(FILE *stream, const char *format, ...);
int sscanf(const char *str, const char *format, ...);
```

## DESCRIPTION

The `scanf()` family of functions scans input according to *format* as described below. This format may contain *conversion specifications*; the results from such conversions, if any, are stored in the locations pointed to by the *pointer* arguments that follow *format*. Each *pointer* argument must be of a type that is appropriate for the value returned by the corresponding conversion specification.

If the number of conversion specifications in *format* exceeds the number of *pointer* arguments, the results are undefined. If the number of *pointer* arguments exceeds the number of conversion specifications, then the excess *pointer* arguments are evaluated, but are otherwise ignored.

The `scanf()` function reads input from the standard input stream *stdin*, `fscanf()` reads input from the stream *pointer* *stream*, and `sscanf()` reads its input from the character string pointed to by *str*.

The *format* string consists of a sequence of *directives* which describe how to process the sequence of input characters. If processing of a directive fails, no further input is read, and `scanf()` returns. A "failure" can be either of the following: *input failure*, meaning that input characters were unavailable, or *matching failure*, meaning that the input was inappropriate (see below).

A directive is one of the following:

- A sequence of white-space characters (space, tab, newline, etc.; see `isspace(3)`). This directive matches any amount of white space, including none, in the input.
- An ordinary character (i.e., one other than white space or '%'). This character must exactly match the next character of input.
- A conversion specification, which commences with a '%' (percent) character. A sequence of characters from the input is converted according to this specification, and the result is placed in the corresponding *pointer* argument. If the next item of input does not match the conversion specification, the conversion fails—this is a *matching failure*.

Each *conversion specification* in *format* begins with either the character '%' or the character sequence "%n\$" (see below for the distinction) followed by:

- An optional '\*' assignment-suppression character: `scanf()` reads input as directed by the conversion specification, but discards the input. No corresponding *pointer* argument is required, and this specification is not included in the count of successful assignments returned by `scanf()`.
- For decimal conversions, an optional quote character (''). This specifies that the input number may include thousands' separators as defined by the `LC_NUMERIC` category of the current locale. (See `setlocale(3)`.) The quote character may precede or follow the '\*' assignment-suppression character.
- An optional decimal integer which specifies the *maximum field width*. Reading of characters stops either when this maximum is reached or when a nonmatching character is found, whichever happens first. Most conversions discard initial white space characters (the exceptions are noted below), and these discarded characters don't count toward the maximum field width. String input conversions store a terminating null byte ('\0') to mark the end of the input; the maximum field width does not include this terminator.
- An optional *type modifier character*. For example, the **I** type modifier is used with integer conversions such as `%d` to specify that the corresponding *pointer* argument refers to a *long int* rather than a pointer to an *int*.

- A *conversion specifier* that specifies the type of input conversion to be performed.

The conversion specifications in *format* are of two forms, either beginning with '%' or beginning with "%n\$". The two forms should not be mixed in the same *format* string, except that a string containing "%n\$" specifications can include %% and %\*. If *format* contains '%' specifications, then these correspond in order with successive *pointer* arguments. In the "%n\$" form (which is specified in POSIX.1-2001, but not C99), *n* is a decimal integer that specifies that the converted input should be placed in the location referred to by the *n*-th *pointer* argument following *format*.

## Conversions

The following *type modifier characters* can appear in a conversion specification:

- I** Indicates either that the conversion will be one of **d**, **i**, **o**, **u**, **x**, **X**, or **n** and the next pointer is a pointer to a *long int* or *unsigned long int* (rather than *int*), or that the conversion will be one of **e**, **f**, or **g** and the next pointer is a pointer to *double* (rather than *float*). Specifying two **I** characters is equivalent to **L**. If used with **%c** or **%s**, the corresponding parameter is considered as a pointer to a wide character or wide-character string respectively.
- L** Indicates that the conversion will be either **e**, **f**, or **g** and the next pointer is a pointer to *long double* or the conversion will be **d**, **i**, **o**, **u**, or **x** and the next pointer is a pointer to *long long*.

The following *conversion specifiers* are available:

- %** Matches a literal '%'. That is, %% in the format string matches a single input '%' character. No conversion is done (but initial white space characters are discarded), and assignment does not occur.
- d** Matches an optionally signed decimal integer; the next pointer must be a pointer to *int*.
- i** Matches an optionally signed integer; the next pointer must be a pointer to *int*. The integer is read in base 16 if it begins with *0x* or *0X*, in base 8 if it begins with *0*, and in base 10 otherwise. Only characters that correspond to the base are used.
- u** Matches an unsigned decimal integer; the next pointer must be a pointer to *unsigned int*.
- x** Matches an unsigned hexadecimal integer; the next pointer must be a pointer to *unsigned int*.
- f** Matches an optionally signed floating-point number; the next pointer must be a pointer to *float*.
- s** Matches a sequence of non-white-space characters; the next pointer must be a pointer to the initial element of a character array that is long enough to hold the input sequence and the terminating null byte ('\0'), which is added automatically. The input string stops at white space or at the maximum field width, whichever occurs first.
- c** Matches a sequence of characters whose length is specified by the *maximum field width* (default 1); the next pointer must be a pointer to *char*, and there must be enough room for all the characters (no terminating null byte is added). The usual skip of leading white space is suppressed. To skip white space first, use an explicit space in the format.
- p** Matches a pointer value (as printed by **%p** in `printf(3)`); the next pointer must be a pointer to a pointer to *void*.

## RETURN VALUE

On success, these functions return the number of input items successfully matched and assigned; this can be fewer than provided for, or even zero, in the event of an early matching failure.

The value **EOF** is returned if the end of input is reached before either the first successful conversion or a matching failure occurs. **EOF** is also returned if a read error occurs, in which case the error indicator for the stream (see `ferror(3)`) is set, and *errno* is set to indicate the error.

sigaction(2)

sigaction(2)

#### NAME

sigaction – POSIX signal handling functions.

#### SYNOPSIS

```
#include <signal.h>
```

```
int sigaction(int signum, const struct sigaction *act, struct sigaction *oldact);
```

#### DESCRIPTION

The **sigaction** system call is used to change the action taken by a process on receipt of a specific signal.

*signum* specifies the signal and can be any valid signal except **SIGKILL** and **SIGSTOP**.

If *act* is non-null, the new action for signal *signum* is installed from *act*. If *oldact* is non-null, the previous action is saved in *oldact*.

The **sigaction** structure is defined as something like

```
struct sigaction {
    void (*sa_handler)(int);
    void (*sa_sigaction)(int, siginfo_t *, void *);
    sigset_t sa_mask;
    int sa_flags;
    void (*sa_restorer)(void);
}
```

On some architectures a union is involved - do not assign to both *sa\_handler* and *sa\_sigaction*.

The *sa\_restorer* element is obsolete and should not be used. POSIX does not specify a *sa\_restorer* element.

*sa\_handler* specifies the action to be associated with *signum* and may be **SIG\_DFL** for the default action, **SIG\_IGN** to ignore this signal, or a pointer to a signal handling function.

*sa\_mask* gives a mask of signals which should be blocked during execution of the signal handler. In addition, the signal which triggered the handler will be blocked, unless the **SA\_NODEFER** or **SA\_NOMASK** flags are used.

*sa\_flags* specifies a set of flags which modify the behaviour of the signal handling process. It is formed by the bitwise OR of zero or more of the following:

##### SA\_NOCLDSTOP

If *signum* is **SIGCHLD**, do not receive notification when child processes stop (i.e., when child processes receive one of **SIGSTOP**, **SIGTSTP**, **SIGTTIN** or **SIGTTOU**).

##### SA\_RESTART

Provide behaviour compatible with BSD signal semantics by making certain system calls restartable across signals.

#### RETURN VALUES

**sigaction** returns 0 on success and -1 on error.

#### ERRORS

##### EINVAL

An invalid signal was specified. This will also be generated if an attempt is made to change the action for **SIGKILL** or **SIGSTOP**, which cannot be caught.

#### SEE ALSO

**kill(1)**, **kill(2)**, **killpg(2)**, **pause(2)**, **sigsetops(3)**,

sigsetops(3C)

sigsetops(3C)

#### NAME

sigsetops, sigemptyset, sigfillset, sigaddset, sigdelset, sigismember – manipulate sets of signals

#### SYNOPSIS

```
#include <signal.h>
```

```
int sigemptyset(sigset_t *set);
```

```
int sigfillset(sigset_t *set);
```

```
int sigaddset(sigset_t *set, int signo);
```

```
int sigdelset(sigset_t *set, int signo);
```

```
int sigismember(sigset_t *set, int signo);
```

#### DESCRIPTION

These functions manipulate *sigset\_t* data types, representing the set of signals supported by the implementation.

**sigemptyset()** initializes the set pointed to by *set* to exclude all signals defined by the system.

**sigfillset()** initializes the set pointed to by *set* to include all signals defined by the system.

**sigaddset()** adds the individual signal specified by the value of *signo* to the set pointed to by *set*.

**sigdelset()** deletes the individual signal specified by the value of *signo* from the set pointed to by *set*.

**sigismember()** checks whether the signal specified by the value of *signo* is a member of the set pointed to by *set*.

Any object of type *sigset\_t* must be initialized by applying either **sigemptyset()** or **sigfillset()** before applying any other operation.

#### RETURN VALUES

Upon successful completion, the **sigismember()** function returns a value of one if the specified signal is a member of the specified set, or a value of 0 if it is not. Upon successful completion, the other functions return a value of 0. Otherwise a value of -1 is returned and **errno** is set to indicate the error.

#### ERRORS

**sigaddset()**, **sigdelset()**, and **sigismember()** will fail if the following is true:

**EINVAL** The value of the *signo* argument is not a valid signal number.

**sigfillset()** will fail if the following is true:

**EFAULT** The *set* argument specifies an invalid address.

#### SEE ALSO

**sigaction(2)**, **sigpending(2)**, **sigprocmask(2)**, **sigsuspend(2)**, **attributes(5)**, **signal(5)**

strcmp(3)

strcmp(3)

**NAME**

strcmp, strncmp – compare two strings

**SYNOPSIS**

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

```
int strcmp(const char *s1, const char *s2);
```

```
int strncmp(const char *s1, const char *s2, size_t n);
```

**DESCRIPTION**

The **strcmp()** function compares the two strings *s1* and *s2*. It returns an integer less than, equal to, or greater than zero if *s1* is found, respectively, to be less than, to match, or be greater than *s2*.

The **strncmp()** function is similar, except it only compares the first (at most) *n* characters of *s1* and *s2*.

**RETURN VALUE**

The **strcmp()** and **strncmp()** functions return an integer less than, equal to, or greater than zero if *s1* (or the first *n* bytes thereof) is found, respectively, to be less than, to match, or be greater than *s2*.

**CONFORMING TO**

SVr4, 4.3BSD, C89, C99.

**SEE ALSO**

**bcmp(3)**, **memcmp(3)**, **strcascmp(3)**, **strcoll(3)**, **strncascmp(3)**, **wscmp(3)**, **wscnmp(3)**

\_\_sync\_fetch\_and\_add/\_\_sync\_fetch\_and\_sub(3)

\_\_sync\_fetch\_and\_add/\_\_sync\_fetch\_and\_sub(3)

**NAME**

\_\_sync\_fetch\_and\_add – Atomic addition/  
\_\_sync\_fetch\_and\_sub – Atomic subtraction

**SYNOPSIS**

```
type __sync_fetch_and_add(type *ptr, type value,...)
```

```
type __sync_fetch_and_sub(type *ptr, type value,...)
```

**DESCRIPTION** **\_\_sync\_fetch\_and\_add/\_\_sync\_fetch\_and\_sub**

These GCC-built-in functions perform the operation suggested by the name, and return the value that had previously been in memory. That is, operations on integer operands have the following semantics:

```
{ tmp = *ptr; *ptr += value; return tmp; }  
{ tmp = *ptr; *ptr -= value; return tmp; }
```

Both **\_\_sync\_fetch\_and\_add()** and **\_\_sync\_fetch\_and\_sub()** are overloaded such that they work on multiple types.

In most cases, these builtins are considered a full barrier. That is, no memory operand will be moved across the operation, either forward or backward. Further, instructions will be issued as necessary to prevent the processor from speculating loads across the operation and from queuing stores after the operation.

**RETURN VALUE**

The value that had previously been in memory pointed to by *ptr*.



triangle(3)

triangle(3)

**NAME**

countPoints – count the number of integer coordinates on the boundary of and inside the triangle

**SYNOPSIS**

```
#include "triangle.h"
```

```
void countPoints(const struct triangle *tri, int* boundary, int* interior);
```

**DESCRIPTION**

Given a triangle *tri* with all corners on integer coordinates (see **struct coordinate**), **countPoints()** counts the number of points (on integer coordinates) on the boundary of the triangle and the number of points inside the triangle.

The parameters *boundary* and *interior* are output parameters that receive the number of points found on the boundary and inside the triangle, respectively.

The **struct coordinate** represents a two-dimensional coordinate in the Cartesian coordinate system. The **struct triangle** stores the three coordinates that make up a triangle.

```
struct coordinate {
```

```
    int x;
```

```
    int y;
```

```
};
```

```
struct triangle {
```

```
    struct coordinate point[3];
```

```
};
```

**RETURN VALUES**

The **countPoints()** function returns no value.