

accept(2)

accept(2)

bind(2)

bind(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, *ms*, 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, *ms*, 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)`

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.

fopen/fdopen/filenop(3)

open/fdopen/filenop(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 **flopen** 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 **flopen** is closed. The result of applying **flopen** 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)**

fork(2)

fork(2)

NAME

fork – create a child process

SYNOPSIS

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

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

DESCRIPTION

fork() creates a new process by duplicating the calling process. The new process, referred to as the *child*, is an exact duplicate of the calling process, referred to as the *parent*, except for the following points:

- * The child has its own unique process ID, and this PID does not match the ID of any existing process group (**sepgid(2)**).
- * The child's parent process ID is the same as the parent's process ID.
- * The child does not inherit its parent's memory locks (**mblock(2)**, **mlockall(2)**).
- * Process resource utilizations (**getusage(2)**) and CPU time counters (**times(2)**) are reset to zero in the child.
- * The child's set of pending signals is initially empty (**sigpending(2)**).
- * The child does not inherit semaphore adjustments from its parent (**semop(2)**).
- * The child does not inherit record locks from its parent (**fcntl(2)**).
- * The child does not inherit timers from its parent (**setitimer(2)**, **alarm(2)**, **timer_create(2)**).
- * The child does not inherit outstanding asynchronous I/O operations from its parent (**aio_read(3)**, **aio_write(3)**), nor does it inherit any asynchronous I/O contexts from its parent (see **io_setup(2)**).

The process attributes in the preceding list are all specified in POSIX.1-2001. The parent and child also differ with respect to the following Linux-specific process attributes:

- * The child does not inherit directory change notifications (dnotify) from its parent (see the description of **F_NOTIFY** in **fcntl(2)**).
 - * The **prctl(2)** **PR_SET_PDEATHSIG** setting is reset so that the child does not receive a signal when its parent terminates.
 - * Memory mappings that have been marked with the **madvisec(2)** **MADV_DONTFORK** flag are not inherited across a **fork()**.
 - * The termination signal of the child is always **SIGCHLD** (see **clone(2)**).
- Note the following further points:
- * The child process is created with a single thread — the one that called **fork()**. The entire virtual address space of the parent is replicated in the child, including the states of mutexes, condition variables, and other pthreads objects; the use of **pthread_atfork(3)** may be helpful for dealing with problems that this can cause.
 - * The child inherits copies of the parent's set of open file descriptors. Each file descriptor in the child refers to the same open file description (see **open(2)**) as the corresponding file descriptor in the parent. This means that the two descriptors share open file status flags, current file offset, and signal-driven I/O attributes (see the description of **F_SETOWN** and **F_SETSIG** in **fcntl(2)**).
 - * The child inherits copies of the parent's set of open message queue descriptors (see **mq_overview(7)**). Each descriptor in the child refers to the same open message queue as the corresponding descriptor in the parent. This means that the two descriptors share the same flags (*mq_flags*).
 - * The child inherits copies of the parent's set of open directory streams (see **opendir(3)**). POSIX.1-2001 says that the corresponding directory streams in the parent and child *may* share the directory stream positioning; on Linux/glibc they do not.

fork(2)

fork(2)

RETURN VALUE

On success, the PID of the child process is returned in the parent, and 0 is returned in the child. On failure, `-1` is returned in the parent, no child process is created, and `errno` is set appropriately.

ERRORS

EAGAIN

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

ENOMEM

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

SVr4, 4.3BSD, POSIX.1-2001.

NOTES

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.

Since version 2.3.3, rather than invoking the kernel's `fork()` system call, the glibc `fork()` wrapper that is provided as part of the NPTL threading implementation invokes `clone(2)` with flags that provide the same effect as the traditional system call. The glibc wrapper invokes any fork handlers that have been established using `pthread_atfork(3)`.

EXAMPLE

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

SEE ALSO

`clone(2)`, `execve(2)`, `setrlimit(2)`, `unshare(2)`, `vfork(2)`, `wait(2)`, `daemon(3)`, `capabilities(7)`, `credentials(7)`

COLOPHON

This page is part of release 3.27 of the Linux *man-pages* project. A description of the project, and information about reporting bugs, can be found at <http://www.kernel.org/doc/man-pages/>.

gets(3)

gets(3)

NAME

`gets`, `fgets` – get a string from a stream
`fputs`, `puts` – output of strings

SYNOPSIS

```
#include <stdio.h>
char *gets(char *s);
char *fgets(char *s, int n, FILE *stream);
int fputs(const char *s, FILE *stream);
int puts(const char *s);
```

DESCRIPTION

The `gets()` function reads characters from the standard input stream (see [intro\(3\)](#)), `stdin`, into the array pointed to by `s`, until a newline character is read or an end-of-file condition is encountered. The newline character is discarded and the string is terminated with a null character.

The `fgets()` function reads characters from the `stream` into the array pointed to by `s`, until `n-1` characters are read, or a newline character is read and transferred to `s`, or an end-of-file condition is encountered. The string is then terminated with a null character.

When using `gets()`, if the length of an input line exceeds the size of `s`, indeterminate behavior may result. For this reason, it is strongly recommended that `gets()` be avoided in favor of `fgets()`.

RETURN VALUES

If end-of-file is encountered and no characters have been read, no characters are transferred to `s` and a null pointer is returned. If a read error occurs, such as trying to use these functions on a file that has not been opened for reading, a null pointer is returned and the error indicator for the stream is set. If end-of-file is encountered, the EOF indicator for the stream is set. Otherwise `s` is returned.

ERRORS

The `gets()` and `fgets()` functions will fail if data needs to be read and:

EOVERFLOW The file is a regular file and an attempt was made to read at or beyond the offset maximum associated with the corresponding `stream`.

DESCRIPTION

`fputs()` writes the string `s` to `stream`, without its trailing `'\0'`.

`puts()` writes the string `s` and a trailing newline to `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

`puts()` and `fputs()` return a non - negative number on success, or **EOF** on error.

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_1    sin6_family; /* AF_INET6 */
    uint6_1    sin6_port; /* port number */
    uint32_1    sin6_flowinfo; /* IPv6 flow information */
    struct in6_addr sin6_addr; /* IPv6 address */
    uint32_1    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

[cmsnsg\(3\)](#), [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\(7\)](#)

malloc(3)

malloc(3)

NAME

calloc, malloc, free, realloc – Allocate and free dynamic memory

SYNOPSIS

```
#include <stdlib.h>
```

```
void *calloc(size_t membr, size_t size);  
void *malloc(size_t size);  
void free(void *ptr);  
void *realloc(void *ptr, size_t size);
```

DESCRIPTION

calloc() allocates memory for an array of *membr* elements of *size* bytes each and returns a pointer to the allocated memory. The memory is set to zero.

malloc() allocates *size* bytes and returns a pointer to the allocated memory. The memory is not cleared.

free() frees the memory space pointed to by *ptr*, which must have been returned by a previous call to **malloc()**, **calloc()** or **realloc()**. Otherwise, or if **free(ptr)** has already been called before, undefined behaviour occurs. If *ptr* is **NULL**, no operation is performed.

realloc() changes the size of the memory block pointed to by *ptr* to *size* bytes. The contents will be unchanged to the minimum of the old and new sizes; newly allocated memory will be uninitialized. If *ptr* is **NULL**, the call is equivalent to **malloc(size)**; if *size* is equal to zero, the call is equivalent to **free(ptr)**. Unless *ptr* is **NULL**, it must have been returned by an earlier call to **malloc()**, **calloc()** or **realloc()**.

RETURN VALUE

For **calloc()** and **malloc()**, the value returned is a pointer to the allocated memory, which is suitably aligned for any kind of variable, or **NULL** if the request fails.

free() returns no value.

realloc() returns a pointer to the newly allocated memory, which is suitably aligned for any kind of variable and may be different from *ptr*, or **NULL** if the request fails. If *size* was equal to 0, either **NULL** or a pointer suitable to be passed to **free()** is returned. If **realloc()** fails the original block is left untouched - it is not freed or moved.

CONFORMING TO

ANSI C

SEE ALSO

brk(2), **posix_memalign(3)**

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_t(3)**

socket(2)

socket(2)

printf(3)

printf(3)

NAME

socket – create an endpoint for communication

SYNOPSIS

```
#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. The currently understood families are:

PF_INET ARPA Internet protocols

The socket has the indicated *type*, which specifies the communication semantics. Currently defined types are:

SOCK_STREAM

SOCK_DGRAM

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).

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 **ETIMEOUT** as the specific code in the global variable **errno**. 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.

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.

ENOMEM Insufficient user memory is available.

SEE ALSO

close(2), **read(2)**, **write(2)**, **accept(3N)**, **bind(3N)**, **connect(3N)**, **listen(3N)**,

NAME

printf, fprintf, sprintf, vsprintf, vprintf, vfprintf, vsprintf – formatted output conversion

SYNOPSIS

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

```
int printf(const char * format, ...);
int fprintf(FILE * stream, const char * format, ...);
int sprintf(char * str, const char * format, ...);
int vsprintf(char * str, size_t size, const char * format, ...);
...

```

DESCRIPTION

The functions in the **printf()** family produce output according to a *format* as described below. The functions **printf()** and **vprintf()** write output to *stdout*, the standard output stream; **fprintf()** and **vfprintf()** write output to the given output stream; **sprintf()**, **vsprintf()**, **vsprintf()** and **vvsprintf()** write to the character string *str*.

The functions **sprintf()** and **vsprintf()** write at most *size* bytes (including the trailing null byte '\0') to *str*.

These eight functions write the output under the control of a *format* string that specifies how subsequent arguments (or arguments accessed via the variable-length argument facilities of **stdarg(3)**) are converted for output.

Return value

Upon successful return, these functions return the number of characters printed (not including the trailing '\0' used to end output to strings).

The functions **sprintf()** and **vsprintf()** do not write more than *size* bytes (including the trailing '\0'). If the output was truncated due to this limit then the return value is the number of characters (not including the trailing '\0') which would have been written to the final string if enough space had been available. Thus, a return value of *size* or more means that the output was truncated. (See also below under NOTES.)

If an output error is encountered, a negative value is returned.

Format of the format string

The format string is a character string, beginning and ending in its initial shift state, if any. The format string is composed of zero or more directives: ordinary characters (not '%'), which are copied unchanged to the output stream, and conversion specifications, each of which results in fetching zero or more subsequent arguments. Each conversion specification is introduced by the character '%', and ends with a *conversion specifier*. In between there may be (in this order) *zero* or more *flags*, an optional *minimum field width*, an optional *precision* and an optional *length modifier*.

The conversion specifier

A character that specifies the type of conversion to be applied. An example for a conversion specifier is:

s The *conv char* * argument is expected to be a pointer to an array of character type (pointer to a string). Characters from the array are written up to (but not including) a terminating null byte ('\0'); if a precision is specified, no more than the number specified are written. If a precision is given, no null byte need be present; if the precision is not specified, or is greater than the size of the array, the array must contain a terminating null byte.

SEE ALSO

printf(1), **asprintf(3)**, **dprintf(3)**, **scanf(3)**, **setlocale(3)**, **wctomb(3)**, **wprintf(3)**, **locale(5)**