

**NAME**

exec, execl, execlp, execlxe, execlxp, execlxpe, execlxpep, execlxpep - execute a file

**SYNOPSIS**

```
#include <unistd.h>

int execl(const char *path, const char *arg0, ..., const char *argn, char * /*NULL*/);

int execlp(const char *path, char *const argv[], ...);

int execlxe(const char *path, char *const argv[], char *const envp[],
char * /*NULL*/, char *const envpl[]);

int execlxpe(const char *path, char *const argv[], char *const envp[]);

int execlxpep(const char *file, const char *arg0, ..., const char *argn, char * /*NULL*/);

int execlxpep(const char *file, char *const argv[]);
```

**DESCRIPTION**

Each of the functions in the **exec** family overlays a new process image on an old process. The new process image is constructed from an ordinary, executable file. This file is either an executable object file, or a file of data for an interpreter. There can be no return from a successful call to one of these functions because the calling process image is overlaid by the new process image.

When a C program is executed, it is called as follows:

```
int main (int argc, char *argv[], char *envp[]);
```

where *argc* is the argument count, *argv* is an array of character pointers to the arguments themselves, and *envp* is an array of character pointers to the environment strings. As indicated, *argc* is at least one, and the first member of the array points to a string containing the name of the file.

The arguments *arg0*, ..., *argn* point to null-terminated character strings. These strings constitute the argument list available to the new process image. Conventionally at least *arg0* should be present. The *arg0* argument points to a string that is the same as *path* (or the last component of *path*). The list of argument strings is terminated by a (**char** \***0**) argument.

The *argv* argument is an array of character pointers to null-terminated strings. These strings constitute the argument list available to the new process image. By convention, *argv* must have at least one member, and it should point to a string that is the same as *path* (or its last component). The *argv* argument is terminated by a null pointer.

The *path* argument points to a path name that identifies the new process file.

The *file* argument points to the new process file. If *file* does not contain a slash character, the path prefix for this file is obtained by a search of the directories passed in the **PATH** environment variable (see **environ(5)**).

File descriptors open in the calling process remain open in the new process.

Signals that are being caught by the calling process are set to the default disposition in the new process image (see **signal(3C)**). Otherwise, the new process image inherits the signal dispositions of the calling process.

**RETURN VALUES**

If a function in the **exec** family returns to the calling process, an error has occurred; the return value is **-1** and **errno** is set to indicate the error.



printf(3) printf(3)

**NAME**  
printf, fprintf, sprintf, snprintf, vprintf, vfprintf, vsprintf, vsnprintf – 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 snprintf(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 function **printf()** writes output to *stdout*, the standard output stream; **fprintf()** writes output to the given output stream; **sprintf()** and **snprintf()**, write to the character string *str*.  
The function **snprintf()** writes at most *size* bytes (including the trailing null byte '\0') to *str*.  
These 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.  
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:  
**o, u, x, X**  
The *unsigned int* argument is converted to unsigned octal (**o**), unsigned decimal (**u**), or unsigned hexadecimal (**x** and **X**) notation.

**s**  
The *const 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)**

qsort(3) qsort(3)

**NAME**  
qsort – sorts an array

**SYNOPSIS**  
#include <stdlib.h>

```
void qsort(void *base, size_t nmemb, size_t size,
           int(*compar)(const void *, const void *));
```

**DESCRIPTION**  
The **qsort()** function sorts an array with *nmemb* elements of size *size*. The *base* argument points to the start of the array.  
The contents of the array are sorted in ascending order according to a comparison function pointed to by *compar*, which is called with two arguments that point to the objects being compared.  
The comparison function must return an integer less than, equal to, or greater than zero if the first argument is considered to be respectively less than, equal to, or greater than the second. If two members compare as equal, their order in the sorted array is undefined.

**RETURN VALUE**  
The **qsort()** function returns no value.

**SEE ALSO**  
**sort(1), alphasort(3), strcmp(3), versionsort(3)**

**ATTRIBUTES**  
**Multithreading (see pthreads(7))**  
The **qsort()** function is thread-safe if the comparison function *compar* does not access any global variables.

**read(2)** (of more than zero bytes). Other routines, like **mmap(2)**, may or may not update *st\_atime*.  
 The field *st\_mtime* is changed by file modifications, for example, by **mknod(2)**, **truncate(2)**, **utime(2)** and **write(2)** (of more than zero bytes). Moreover, *st\_mtime* of a directory is changed by the creation or deletion of files in that directory. The *st\_mtime* field is *not* changed for changes in owner, group, hard link count, or mode.

The field *st\_ctime* is changed by writing or by setting inode information (i.e., owner, group, link count, mode, etc.).

The following POSIX macros are defined to check the file type using the *st\_mode* field:

**S\_ISREG(m)** is it a regular file?  
**S\_ISDIR(m)** directory?  
**S\_ISCHR(m)** character device?  
**S\_ISBLK(m)** block device?  
**S\_ISFIFO(m)** FIFO (named pipe)?  
**S\_ISLNK(m)** symbolic link? (Not in POSIX.1-1996.)  
**S\_ISSOCK(m)** socket? (Not in POSIX.1-1996.)

The following mask values are defined for the file mode component of the *st\_mode* field:

**S\_IRWXU** 00700 owner has read, write, and execute permission  
**S\_IRUSR** 00400 owner has read permission  
**S\_IWUSR** 00200 owner has write permission  
**S\_IXUSR** 00100 owner has execute permission

#### RETURN VALUE

On success, zero is returned. On error,  $-1$  is returned, and *errno* is set appropriately.

#### ERRORS

##### EACCES

Search permission is denied for one of the directories in the path prefix of *path*. (See also **path\_resolution(7)**.)

##### EBADF

*fd* is bad.

##### EFAULT

Bad address.

##### ELOOP

Too many symbolic links encountered while traversing the path.

##### ENAMETOOLONG

File name too long.

##### ENOENT

A component of the path *path* does not exist, or the path is an empty string.

##### ENOMEM

Out of memory (i.e., kernel memory).

##### ENOTDIR

A component of the path is not a directory.

#### SEE ALSO

**access(2)**, **chmod(2)**, **chown(2)**, **fstatat(2)**, **readlink(2)**, **utime(2)**, **capabilities(7)**, **symlink(7)**

**NAME**  
 stat, fstat, lstat — get file status

**SYNOPSIS**  
**#include <sys/types.h>**  
**#include <sys/stat.h>**  
**#include <unistd.h>**

```
int stat(const char *path, struct stat *buf);
int fstat(int fd, struct stat *buf);
int lstat(const char *path, struct stat *buf);
```

#### DESCRIPTION

These functions return information about a file. No permissions are required on the file itself, but — in the case of **stat()** and **lstat()** — execute (search) permission is required on all of the directories in *path* that lead to the file.

**stat()** stats the file pointed to by *path* and fills in *buf*.

**lstat()** is identical to **stat()**, except that if *path* is a symbolic link, then the link itself is stat-ed, not the file that it refers to.

**fstat()** is identical to **lstat()**, except that the file to be stat-ed is specified by the file descriptor *fd*.

All of these system calls return a *stat* structure, which contains the following fields:

```
struct stat {
    dev_t   st_dev;      /* ID of device containing file */
    ino_t   st_ino;     /* inode number */
    mode_t  st_mode;    /* protection */
    nlink_t st_nlink;   /* number of hard links */
    uid_t   st_uid;    /* user ID of owner */
    gid_t   st_gid;    /* group ID of owner */
    dev_t   st_rdev;   /* device ID (if special file) */
    off_t   st_size;   /* total size, in bytes */
    blksize_t st_blksize; /* blocksize for file system I/O */
    blkcnt_t st_blocks; /* number of blocks allocated */
    time_t  st_atime;  /* time of last access */
    time_t  st_mtime;  /* time of last modification */
    time_t  st_ctime;  /* time of last status change */
};
```

The *st\_dev* field describes the device on which this file resides.

The *st\_rdev* field describes the device that this file (inode) represents.

The *st\_size* field gives the size of the file (if it is a regular file or a symbolic link) in bytes. The size of a symlink is the length of the pathname it contains, without a trailing null byte.

The *st\_blocks* field indicates the number of blocks allocated to the file, 512-byte units. (This may be smaller than *st\_size/512* when the file has holes.)

The *st\_blksize* field gives the "preferred" blocksize for efficient file system I/O. (Writing to a file in smaller chunks may cause an inefficient read-modify-rewrite.)

Not all of the Linux file systems implement all of the time fields. Some file system types allow mounting in such a way that file accesses do not cause an update of the *st\_atime* field. (See "noatime" in **mount(8)**.)

The field *st\_atime* is changed by file accesses, for example, by **execve(2)**, **pipe(2)**, **utime(2)** and

string(3)		string(3)	time(2)	time(2)
NAME	strcat, strchr, streamp, strcpy, strdup, strlen, strcat, strcmp, strncmp, strncpy, strstr, strtok – string operations		NAME	time – get time in seconds
SYNOPSIS	<pre>#include &lt;string.h&gt; char *strcat(char *dest, const char *src); char *strcpy(const char *s, int c); int strcmp(const char *s1, const char *s2); char *strchr(char *dest, const char *src); char *strdup(const char *s); size_t strlen(const char *s); char *strncat(char *dest, const char *src, size_t n); int strncmp(const char *s1, const char *s2, size_t n); char *strncpy(char *dest, const char *src, size_t n); char *strrchr(const char *haystack, const char *needle); char *strtok(char *s, const char *delim);</pre>		SYNOPSIS	#include <time.h> time_t time(time_t *tloc);
DESCRIPTION	<p>The string functions perform operations on null-terminated strings.</p>		DESCRIPTION	time() returns the time as the number of seconds since the Epoch, 1970-01-01 00:00:00 +0000 (UTC). If tloc is non-NULL, the return value is also stored in the memory pointed to by tloc. When tloc is NULL, the call cannot fail. RETURN VALUE On success, the value of time in seconds since the Epoch is returned. On error, ((time_t) -1) is returned, and errno is set appropriately.

**NAME** waitpid – wait for child process to change state

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/wait.h>
pid_t waitpid(pid_t pid, int *stat_loc, int options);
```

**DESCRIPTION**

`waitpid()` suspends the calling process until one of its children changes state; if a child process changed state prior to the call to `waitpid()`, return is immediate. `pid` specifies a set of child processes for which status is requested.

If `pid` is equal to `(pid_t)-1`, status is requested for any child process.

If `pid` is greater than `(pid_t)0`, it specifies the process ID of the child process for which status is requested.

If `pid` is equal to `(pid_t)0` status is requested for any child process whose process group ID is equal to that of the calling process.

If `pid` is less than `(pid_t)-1`, status is requested for any child process whose process group ID is equal to the absolute value of `pid`.

If `waitpid()` returns because the status of a child process is available, then that status may be evaluated with the macros defined by `wstat(5)`. If the calling process had specified a non-zero value of `stat_loc`, the status of the child process will be stored in the location pointed to by `stat_loc`.

The `options` argument is constructed from the bitwise inclusive OR of zero or more of the following flags, defined in the header `<sys/wait.h>`:

**WCONTINUED** The status of any continued child process specified by `pid`, whose status has not been reported since it continued, is also reported to the calling process.

**WNOHANG** `waitpid()` will not suspend execution of the calling process if status is not immediately available for one of the child processes specified by `pid`.

**WNOWAIT** Keep the process whose status is returned in `stat_loc` in a waitable state. The process may be waited for again with identical results.

If `wstatus` is not NULL, `wait()` and `waitpid()` store status information in the `int` to which it points. This integer can be inspected with the following macros (which take the integer itself as an argument, not a pointer to it, as is done in `wait()` and `waitpid()`):

**WIFEXITED(*wstatus*)** returns true if the child terminated normally, that is, by calling `exit(3)` or `_exit(2)`, or by returning from `main()`.

**WEXITSTATUS(*wstatus*)** returns the exit status of the child. This consists of the least significant 8 bits of the `status` argument that the child specified in a call to `exit(3)` or `_exit(2)` or as the argument for a return statement in `main()`. This macro should be employed only if **WIFEXITED** returned true.

**WIFSIGNALED(*wstatus*)** returns true if the child process was terminated by a signal.

**WTERMSIG(*wstatus*)** returns the number of the signal that caused the child process to terminate. This macro should be employed only if **WIFSIGNALED** returned true.

**RETURN VALUES**

If `waitpid()` returns because the status of a child process is available, this function returns a value equal to the process ID of the child process for which status is reported. If `waitpid()` returns due to the delivery of a signal to the calling process, `-1` is returned and `errno` is set to `EINTR`. If this function was invoked with

**WNOHANG** set in `options`, it has at least one child process specified by `pid` for which status is not available, and status is not available for any process specified by `pid`, `0` is returned. Otherwise, `-1` is returned, and `errno` is set to indicate the error.

**ERRORS**

`waitpid()` will fail if one or more of the following is true:

**ECHILD** The process or process group specified by `pid` does not exist or is not a child of the calling process or can never be in the states specified by `options`.

**EINTR** `waitpid()` was interrupted due to the receipt of a signal sent by the calling process.

**EINVAL** An invalid value was specified for `options`.

**SEE ALSO**

`exec(2)`, `exit(2)`, `fork(2)`, `sigaction(2)`, `wstat(5)`