

exec(2)

exec(2)

**NAME**

exec, execl, execlx, execlx, execlp, execlpx – execute a file

**SYNOPSIS**

```
#include <unistd.h>

int execl(const char *path, const char *arg0, ..., const char *argn, char */*NULL*/);
int execlx(const char *path, char *const argv[]);
int execlx(const char *path, char *const argv[], ..., const char *argn, char */*NULL*/);
int execl(const char *path, char *const arg0[, ..., const char *argn, char */*NULL*/];
char */*NULL*/[, char *const envp[]]);
int execl(const char *path, char *const argv[] char *const envp[]);
int execlp(const char *file, const char *arg0, ..., const char *argn, char */*NULL*/);
int execlpx(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.

stat(2)

stat(2)

**NAME**

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

**DESCRIPTION**

These functions return information about the specified file. You do not need any access rights to the file to get this information but you need search rights to all directories named in the path leading to the file. **stat** stats the file pointed to by *path* and fills in *buf*.

**lstat** is identical to **stat**, except in the case of a symbolic link, where the link itself is stat-ed, not the file that it refers to.

They all return a *stat* structure, which contains the following fields:

```
struct stat {
    dev_t      st_dev; /* device */
    ino_t      st_ino; /* inode */
    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 type (if inode device) */
    off_t      st_size; /* total size, in bytes */
    blksize_t  st_blksize; /* blocksize for filesystem 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 value *st\_size* gives the size of the file (if it is a regular file or a symlink) in bytes. The size of a symlink is the length of the pathname it contains, without trailing NUL.

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

```
S_ISREG(m)    is it a regular file?
S_ISDIR(m)    directory?
```

**RETURN VALUE**

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

**ERRORS**

- EACCESS** Search permission is denied for one of the directories in the path prefix of *path*.
- ENOENT** A component of *path* does not exist, or *path* is an empty string.
- ENOTDIR** A component of the path prefix of *path* is not a directory.

wait(2)

wait(2)

**NAME**

wait, waitpid, waitid – wait for process to change state

**SYNOPSIS**

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

```
pid_t wait(int *status);
```

```
pid_t waitpid(pid_t pid, int *status, int options);
```

**DESCRIPTION**

All of these system calls are used to wait for state changes in a child of the calling process, and obtain information about the child whose state has changed. A state change is considered to be: the child terminated; the child was stopped by a signal; or the child was resumed by a signal. In the case of a terminated child, performing a wait allows the system to release the resources associated with the child; if a wait is not performed, then the terminated child remains in a "zombie" state (see NOTES below).

If a child has already changed state, then these calls return immediately.

The **wait()** system call suspends execution of the calling process until one of its children terminates. The **waitpid()** system call suspends execution of the calling process until a child specified by *pid* argument has changed state. By default, **waitpid()** waits only for terminated children, but this behavior is modifiable via the *options* argument, as described below.

The value of *pid* can be:

- < -1 meaning wait for any child process whose process group ID is equal to the absolute value of *pid*.
- 1 meaning wait for any child process.
- 0 meaning wait for any child process whose process group ID is equal to that of the calling process.
- > 0 meaning wait for the child whose process ID is equal to the value of *pid*.

The value of *options* is an OR of zero or more of the following constants:

**WNOHANG** return immediately if no child has exited.

If *status* 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(status)**

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

**WEXITSTATUS(status)**

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 only be employed if **WIFEXITED** returned true.

**RETURN VALUE**

**wait()**: on success, returns the process ID of the terminated child; on error, -1 is returned.

**waitpid()**: on success, returns the process ID of the child whose state has changed; if **WNOHANG** was specified and one or more child(ren) specified by *pid* exist, but have not yet changed state, then 0 is returned. On error, -1 is returned.