

NAME

exec, execl, execv, execl, execve, execlp, execvp – execute a file

SYNOPSIS

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

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

```
int execv(const char *path, char *const argv[]);
```

```
int execl(const char *path, char *const arg0[], ..., const char *argn,
char * /*NULL*/, char *const envp[]);
```

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

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

```
int execvp(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.

NAME

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

SYNOPSIS

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

```
void *calloc(size_t nmemb, 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 *nmemb* 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)**

memset(3)

memset(3)

NAME

memset – fill memory with a constant byte

SYNOPSIS

#include <string.h>

void *memset(void *s, int c, size_t n);

DESCRIPTION

The **memset()** function fills the first *n* bytes of the memory area pointed to by *s* with the constant byte *c*.

RETURN VALUE

The **memset()** function returns a pointer to the memory area *s*.

CONFORMING TO

POSIX.1-2001, POSIX.1-2008, C89, C99, SVr4, 4.3BSD.

SEE ALSO

bstring(3), bzero(3), swab(3), wmemset(3)

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 **snprintf()** and **vsnprintf()** 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), wcrctomb(3), wprintf(3), locale(5)

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);
```

Feature Test Macro Requirements for glibc (see **feature_test_macros(7)**):

```
lstat(): _BSD_SOURCE || _XOPEN_SOURCE >= 500
```

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 **stat()**, 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)**, **mknod(2)**, **pipe(2)**, **utime(2)** and **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.)

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)**

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)**, **strcasemp(3)**, **strcoll(3)**, **strncasemp(3)**, **wscmp(3)**, **wscncmp(3)**

strtok(3)

strtok(3)

NAME

strtok, strtok_r – extract tokens from strings

SYNOPSIS

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

```
char *strtok(char *str, const char *delim);
```

```
char *strtok_r(char *str, const char *delim, char **saveptr);
```

DESCRIPTION

The **strtok()** function breaks a string into a sequence of zero or more nonempty tokens. On the first call to **strtok()** the string to be parsed should be specified in *str*. In each subsequent call that should parse the same string, *str* must be NULL.

The *delim* argument specifies a set of bytes that delimit the tokens in the parsed string. The caller may specify different strings in *delim* in successive calls that parse the same string.

Each call to **strtok()** returns a pointer to a null-terminated string containing the next token. This string does not include the delimiting byte. If no more tokens are found, **strtok()** returns NULL.

A sequence of calls to **strtok()** that operate on the same string maintains a pointer that determines the point from which to start searching for the next token. The first call to **strtok()** sets this pointer to point to the first byte of the string. The start of the next token is determined by scanning forward for the next nondelimiter byte in *str*. If such a byte is found, it is taken as the start of the next token. If no such byte is found, then there are no more tokens, and **strtok()** returns NULL. (A string that is empty or that contains only delimiters will thus cause **strtok()** to return NULL on the first call.)

The end of each token is found by scanning forward until either the next delimiter byte is found or until the terminating null byte (`\0`) is encountered. If a delimiter byte is found, it is overwritten with a null byte to terminate the current token, and **strtok()** saves a pointer to the following byte; that pointer will be used as the starting point when searching for the next token. In this case, **strtok()** returns a pointer to the start of the found token.

From the above description, it follows that a sequence of two or more contiguous delimiter bytes in the parsed string is considered to be a single delimiter, and that delimiter bytes at the start or end of the string are ignored. Put another way: the tokens returned by **strtok()** are always nonempty strings. Thus, for example, given the string `"aaa;bbb,"`, successive calls to **strtok()** that specify the delimiter string `";,"` would return the strings `"aaa"` and `"bbb"`, and then a null pointer.

The **strtok_r()** function is a reentrant version **strtok()**. The *saveptr* argument is a pointer to a `char *` variable that is used internally by **strtok_r()** in order to maintain context between successive calls that parse the same string. On the first call to **strtok_r()**, *str* should point to the string to be parsed, and the value of *saveptr* is ignored. In subsequent calls, *str* should be NULL, and *saveptr* should be unchanged since the previous call.

Different strings may be parsed concurrently using sequences of calls to **strtok_r()** that specify different *saveptr* arguments.

RETURN VALUE

strtok() and **strtok_r()** return a pointer to the next token, or NULL if there are no more tokens.

ATTRIBUTES

Multithreading (see pthreads(7))

The **strtok()** function is not thread-safe, the **strtok_r()** function is thread-safe.

NAME

strtol, – convert a string to a long integer

SYNOPSIS

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

```
long int strtol(const char *nptr, char **endptr, int base);
```

DESCRIPTION

The **strtol()** function converts the initial part of the string in *nptr* to a long integer value according to the given *base*, which must be between 2 and 36 inclusive, or be the special value 0.

The string may begin with an arbitrary amount of white space (as determined by **isspace(3)**) followed by a single optional '+' or '-' sign. If *base* is zero or 16, the string may then include a "0x" or "0X" prefix, and the number will be read in base 16; otherwise, a zero *base* is taken as 10 (decimal) unless the next character is '0', in which case it is taken as 8 (octal).

The remainder of the string is converted to a *long int* value in the obvious manner, stopping at the first character which is not a valid digit in the given base. (In bases above 10, the letter 'A' in either uppercase or lowercase represents 10, 'B' represents 11, and so forth, with 'Z' representing 35.)

If *endptr* is not NULL, **strtol()** stores the address of the first invalid character in **endptr*. If there were no digits at all, **strtol()** stores the original value of *nptr* in **endptr* (and returns 0). In particular, if *nptr* is not '\0' but **endptr* is '\0' on return, the entire string is valid.

RETURN VALUE

The **strtol()** function returns the result of the conversion, unless the value would underflow or overflow. If an underflow occurs, **strtol()** returns **LONG_MIN**. If an overflow occurs, **strtol()** returns **LONG_MAX**. In both cases, *errno* is set to **ERANGE**.

ERRORS**EINVAL**

(not in C99) The given *base* contains an unsupported value.

ERANGE

The resulting value was out of range.

The implementation may also set *errno* to **EINVAL** in case no conversion was performed (no digits seen, and 0 returned).

NOTES

Since **strtol()** can legitimately return 0, **LONG_MAX**, or **LONG_MIN** on both success and failure, the calling program should set *errno* to 0 before the call, and then determine if an error occurred by checking whether *errno* has a nonzero value after the call.

EXAMPLE

The program shown below demonstrates the use of **strtol()**. The first command-line argument specifies a string from which **strtol()** should parse a number.

Program source

```
#include <stdlib.h>
#include <limits.h>
#include <stdio.h>
#include <errno.h>

int main(int argc, char *argv[])
{
    if (argc < 2) {
        fprintf(stderr, "Usage: %s str\n", argv[0]);
        exit(EXIT_FAILURE);
    }

    errno = 0; /* To distinguish success/failure after call */
    char *endptr;
    long val = strtol(argv[1], &endptr, 10);

    /* Check for various possible errors */
    if ((errno == ERANGE && (val == LONG_MAX || val == LONG_MIN))
        || (errno != 0 && val == 0)) {
        perror("strtol");
        exit(EXIT_FAILURE);
    }

    if (*endptr != '\0')
        printf("Further characters after number: %s\n", endptr);

    printf("strtol() returned %ld\n", val);
}
```

TIME(2) TIME(2)

NAME

time – get time in seconds

SYNOPSIS

```
#include <time.h>
```

```
time_t time(time_t *tloc);
```

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*.

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.

ERRORS

EFAULT

tloc points outside your accessible address space (but see BUGS).

On systems where the C library **time()** wrapper function invokes an implementation provided by the **vdso(7)** (so that there is no trap into the kernel), an invalid address may instead trigger a **SIGSEGV** signal.

NOTES

POSIX.1 defines *seconds since the Epoch* using a formula that approximates the number of seconds between a specified time and the Epoch. This formula takes account of the facts that all years that are evenly divisible by 4 are leap years, but years that are evenly divisible by 100 are not leap years unless they are also evenly divisible by 400, in which case they are leap years. This value is not the same as the actual number of seconds between the time and the Epoch, because of leap seconds and because system clocks are not required to be synchronized to a standard reference. The intention is that the interpretation of seconds since the Epoch values be consistent; see POSIX.1-2008 Rationale A.4.15 for further rationale.

On Linux, a call to **time()** with *tloc* specified as NULL cannot fail with the error **EOverflow**, even on ABIs where *time_t* is a signed 32-bit integer and the clock ticks past the time 2**31 (2038-01-19 03:14:08 UTC, ignoring leap seconds). (POSIX.1 permits, but does not require, the **EOverflow** error in the case where the seconds since the Epoch will not fit in *time_t*.) Instead, the behavior on Linux is undefined when the system time is out of the *time_t* range. Applications intended to run after 2038 should use ABIs with *time_t* wider than 32 bits.

BUGS

Error returns from this system call are indistinguishable from successful reports that the time is a few seconds *before* the Epoch, so the C library wrapper function never sets *errno* as a result of this call.

The *tloc* argument is obsolescent and should always be NULL in new code. When *tloc* is NULL, the call cannot fail.

waitpid(2) waitpid(2)

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*.

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

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)**