

fork(2)

fork – create a child process

NAME

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 (**setpgid(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 (**mlock(2)**, **mlockall(2)**).
- * Process resource utilizations (**getrusage(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 **madvise(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 description 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.

exec(2)

exec, execl, execlp, execlpe, execlpep, execlpevp, execlpevp – execute a file

NAME

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 execlpe(const char *path, char *const arg0[], ..., const char *argn, char * /*NULL*/);
int execlpep(const char *path, char *const argv[] | char *const envp[]);
int execlpevp(const char *file, const char *arg0, ..., const char *argn, char * /*NULL*/);
int execlpevp(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.

fork(2)	<p>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 <i>errno</i> is set appropriately.</p> <p>ERRORS EAGAIN <code>fork()</code> cannot allocate sufficient memory to copy the parent's page tables and allocate a task structure for the child. EAGAIN It was not possible to create a new process because the caller's <code>RLIMIT_NPROC</code> resource limit was encountered. To exceed this limit, the process must have either the <code>CAP_SYS_ADMIN</code> or the <code>CAP_SYS_RESOURCE</code> capability. ENOMEM <code>fork()</code> failed to allocate the necessary kernel structures because memory is tight.</p> <p>CONFORMING TO SVr4, 4.3BSD, POSIX.1-2001.</p> <p>NOTES Under Linux, <code>fork()</code> 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 <code>fork()</code> system call, the glibc <code>fork()</code> wrapper that is provided as part of the NPTL threading implementation invokes <code>clone(2)</code> with flags that provide the same effect as the traditional system call. The glibc wrapper invokes any fork handlers that have been established using <code>pthread_atfork(3)</code>.</p> <p>EXAMPLE See <code>pipe(2)</code> and <code>wait(2)</code>.</p> <p>SEE ALSO <code>clone(2)</code>, <code>execve(2)</code>, <code>setrlimit(2)</code>, <code>vfork(2)</code>, <code>wait(2)</code>, <code>daemon(3)</code>, <code>capabilities(7)</code>, <code>credentials(7)</code></p> <p>COLOPHON This page is part of release 3.27 of the Linux <i>man-pages</i> project. A description of the project, and information about reporting bugs, can be found at http://www.kernel.org/doc/man-pages/.</p>	fork(2)	GSP-Klausur Manual-Auszug	2	2021-07-20	1
opendir/readdir(3)	<p>NAME opendir – open a directory / readdir – read a directory</p> <p>SYNOPSIS <code>#include <sys/types.h></code> <code>#include <dirent.h></code></p> <p><code>DIR *opendir(const char *name);</code> <code>int closedir(DIR *dir);</code> <code>struct dirent *readdir(DIR *dir);</code></p> <p>DESCRIPTION <code>opendir</code> The <code>opendir()</code> function opens a directory stream corresponding to the directory <i>name</i>, and returns a pointer to the directory stream. The stream is positioned at the first entry in the directory.</p> <p>RETURN VALUE The <code>opendir()</code> function returns a pointer to the directory stream. On error, NULL is returned, and <i>errno</i> is set appropriately.</p> <p>DESCRIPTION <code>closedir</code> The <code>closedir()</code> function closes the directory stream associated with <i>dirp</i>. A successful call to <code>closedir()</code> also closes the underlying file descriptor associated with <i>dirp</i>. The directory stream descriptor <i>dirp</i> is not available after this call.</p> <p>RETURN VALUE The <code>closedir()</code> function returns 0 on success. On error, -1 is returned, and <i>errno</i> is set appropriately.</p> <p>DESCRIPTION <code>readdir</code> The <code>readdir()</code> function returns a pointer to a dirent structure representing the next directory entry in the directory stream pointed to by <i>dir</i>. It returns NULL on reaching the end-of-file or if an error occurred. It is safe to use <code>readdir()</code> inside threads if the pointers passed as <i>dir</i> are created by distinct calls to <code>opendir()</code>. The data returned by <code>readdir()</code> is overwritten by subsequent calls to <code>readdir()</code> for the same directory stream. The <i>dirent</i> structure is defined as follows:</p> <pre>struct dirent { long d_ino; /* inode number */ char d_name[256]; /* filename */ };</pre> <p>RETURN VALUE On success, <code>readdir()</code> returns a pointer to a <i>dirent</i> structure. (This structure may be statically allocated; do not attempt to <code>free(3)</code> it.) If the end of the directory stream is reached, NULL is returned and <i>errno</i> is not changed. If an error occurs, NULL is returned and <i>errno</i> is set appropriately. To distinguish end of stream and from an error, set <i>errno</i> to zero before calling <code>readdir()</code> and then check the value of <i>errno</i> if NULL is returned.</p> <p>ERRORS EACCES Permission denied. ENOENT Directory does not exist, or <i>name</i> is an empty string. ENOTDIR <i>name</i> is not a directory.</p>	opendir/readdir(3)	GSP-Klausur Manual-Auszug	3	2021-07-20	1

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 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.

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.)

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.)

The `<sys/stat.h>` header shall define the following symbolic constants for the file mode bits encoded in type `mode_t`, with the indicated numeric values. These macros shall expand to an expression which has a type that allows them to be used, either singly or OR'ed together, as the third argument to `open()` without the need for a `mode_t` cast. The values shall be suitable for use in `#if` preprocessing directives.

Name	Numeric Value	Description
<code>S_IRWXU</code>	0700	Read, write, execute/search by owner.
<code>S_IRUSR</code>	0400	Read permission, owner.
<code>S_IWUSR</code>	0200	Write permission, owner.
<code>S_IXUSR</code>	0100	Execute/search permission, owner.
<code>S_IRWXG</code>	070	Read, write, execute/search by group.
<code>S_IRGRP</code>	040	Read permission, group.
<code>S_IWGRP</code>	020	Write permission, group.
<code>S_IXGRP</code>	010	Execute/search permission, group.
<code>S_IRWXO</code>	07	Read, write, execute/search by others.
<code>S_IROTH</code>	04	Read permission, others.
<code>S_IWOTH</code>	02	Write permission, others.
<code>S_IXOTH</code>	01	Execute/search permission, others.

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)`.)

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)`

string(3)

string(3)

NAME

strcat, strchr, strcmp, strcpy, strdup, strlen, strcat, strcmp, strcpy, strstr, strtok – string operations

SYNOPSIS

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

```
char *strcat(char *dest, const char *src);
```

Append the string *src* to the string *dest*, returning a pointer *dest*.

```
char *strchr(const char *s, int c);
```

Return a pointer to the first occurrence of the character *c* in the string *s*.

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

Compare the strings *s1* with *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*.

```
char *strcpy(char *dest, const char *src);
```

Copy the string *src* to *dest*, returning a pointer to the start of *dest*.

```
char *strdup(const char *s);
```

Return a duplicate of the string *s* in memory allocated using `malloc(3)`.

```
size_t strlen(const char *s);
```

Return the length of the string *s*.

```
char *strncat(char *dest, const char *src, size_t n);
```

Append at most *n* characters from the string *src* to the string *dest*, returning a pointer to *dest*.

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

Compare at most *n* bytes of the 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*.

```
char *strncpy(char *dest, const char *src, size_t n);
```

Copy at most *n* bytes from string *src* to *dest*, returning a pointer to the start of *dest*.

```
char *strstr(const char *haystack, const char *needle);
```

Find the first occurrence of the substring *needle* in the string *haystack*, returning a pointer to the found substring.

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

Extract tokens from the string *s* that are delimited by one of the bytes in *delim*.

DESCRIPTION

The string functions perform operations on null-terminated strings.