

<p>flush(3)</p> <p><b>NAME</b> flush – flush a stream</p> <p><b>SYNOPSIS</b> #include &lt;stdio.h&gt; int fflush(FILE *stream);</p> <p><b>DESCRIPTION</b> For output streams, <b>fflush()</b> forces a write of all user-space buffered data for the given output or update stream via the stream's underlying write function. For input streams associated with seekable files (e.g., disk files, but not pipes or terminals), <b>fflush()</b> discards any buffered data that has been fetched from the underlying file, but has not been consumed by the application. The open status of the stream is unaffected. If the stream argument is NULL, <b>fflush()</b> flushes all open output streams. For a nonlocking counterpart, see <b>unlocked_stdio(3)</b>.</p> <p><b>RETURN VALUE</b> Upon successful completion 0 is returned. Otherwise, <b>EOF</b> is returned and <i>errno</i> is set to indicate the error.</p> <p><b>ERRORS</b> <b>EBADF</b> <i>stream</i> is not an open stream, or is not open for writing. The function <b>fflush()</b> may also fail and set <i>errno</i> for any of the errors specified for <b>write(2)</b>.</p> <p><b>SEE ALSO</b> <b>sync(2)</b>, <b>write(2)</b>, <b>write(3)</b>, <b>fileno(3)</b>, <b>fopen(3)</b>, <b>setbuf(3)</b>, <b>unlocked_stdio(3)</b></p>	<p>flush(3)</p> <p><b>NAME</b> calloc, malloc, free, realloc – Allocate and free dynamic memory</p> <p><b>SYNOPSIS</b> #include &lt;stdlib.h&gt; void *calloc(size_t mmemb, size_t size); void *malloc(size_t size); void free(void *ptr); void *realloc(void *ptr, size_t size);</p> <p><b>DESCRIPTION</b> <b>calloc()</b> allocates memory for an array of <i>mmemb</i> elements of <i>size</i> bytes each and returns a pointer to the allocated memory. The memory is set to zero. <b>malloc()</b> allocates <i>size</i> bytes and returns a pointer to the allocated memory. The memory is not cleared. <b>free()</b> frees the memory space pointed to by <i>ptr</i>, which must have been returned by a previous call to <b>malloc()</b>, <b>calloc()</b> or <b>realloc()</b>. Otherwise, or if <b>free(ptr)</b> has already been called before, undefined behaviour occurs. If <i>ptr</i> is NULL, no operation is performed. <b>realloc()</b> changes the size of the memory block pointed to by <i>ptr</i> to <i>size</i> bytes. The contents will be unchanged to the minimum of the old and new sizes; newly allocated memory will be uninitialized. If <i>ptr</i> is NULL, the call is equivalent to <b>malloc(size)</b>; if <i>size</i> is equal to zero, the call is equivalent to <b>free(ptr)</b>. Unless <i>ptr</i> is NULL, it must have been returned by an earlier call to <b>malloc()</b>, <b>calloc()</b> or <b>realloc()</b>.</p> <p><b>RETURN VALUE</b> For <b>calloc()</b> and <b>malloc()</b>, 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. <b>free()</b> returns no value.</p> <p><b>realloc()</b> returns a pointer to the newly allocated memory, which is suitably aligned for any kind of variable and may be different from <i>ptr</i>, or NULL if the request fails. If <i>size</i> was equal to 0, either NULL or a pointer suitable to be passed to <b>free()</b> is returned. If <b>realloc()</b> fails the original block is left untouched - it is not freed or moved.</p> <p><b>CONFORMING TO</b> ANSI-C</p> <p><b>SEE ALSO</b> <b>brk(2)</b>, <b>posix_memalign(3)</b></p>	<p>malloc(3)</p> <p>malloc(3)</p>	<p>flush(3)</p> <p>flush(3)</p>
<p>GSP-Klausur Manual-Auszug</p> <p>2020-08-11</p> <p>1</p>	<p>GSP-Klausur Manual-Auszug</p> <p>2020-08-11</p> <p>1</p>	<p>GSP-Klausur Manual-Auszug</p> <p>2020-08-11</p> <p>1</p>	<p>GSP-Klausur Manual-Auszug</p> <p>2020-08-11</p> <p>1</p>

<p>opendir/readdir(3)</p> <p>opendir – open a directory / readdir – read a directory</p> <p><b>NAME</b></p> <p><b>SYNOPSIS</b></p> <pre>#include &lt;sys/types.h&gt; #include &lt;dirent.h&gt;  DIR *opendir(const char *name); int closedir(DIR *dir); struct dirent *readdir(DIR *dir);</pre> <p><b>DESCRIPTION</b> <i>opendir</i></p> <p>The <i>opendir</i>() 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><b>RETURN VALUE</b></p> <p>The <i>opendir</i>() function returns a pointer to the directory stream. On error, NULL is returned, and <i>errno</i> is set appropriately.</p> <p><b>DESCRIPTION</b> <i>closedir</i></p> <p>The <i>closedir</i>() function closes the directory stream associated with <i>dirp</i>. A successful call to <i>closedir</i>() 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><b>RETURN VALUE</b></p> <p>The <i>closedir</i>() function returns 0 on success. On error, -1 is returned, and <i>errno</i> is set appropriately.</p> <p><b>DESCRIPTION</b> <i>readdir</i></p> <p>The <i>readdir</i>() 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 <i>readdir</i>() inside threads if the pointers passed as <i>dir</i> are created by distinct calls to <i>opendir</i>().</p> <p>The data returned by <i>readdir</i>() is overwritten by subsequent calls to <i>readdir</i>() for the same directory stream.</p> <p>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><b>RETURN VALUE</b></p> <p>On success, <i>readdir</i>() returns a pointer to a <i>dirent</i> structure. (This structure may be statically allocated; do not attempt to <i>free</i>(3) it.)</p> <p>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 <i>readdir</i>() and then check the value of <i>errno</i> if NULL is returned.</p> <p><b>ERRORS</b></p> <p><b>EACCES</b> Permission denied.</p> <p><b>ENOENT</b> Directory does not exist, or <i>name</i> is an empty string.</p> <p><b>ENOTDIR</b> <i>name</i> is not a directory.</p>	<p>printf/sprintf(3)</p> <p>printf, fprintf, sprintf, snprintf, vprintf, vsprintf, vsnprintf – formatted output conversion</p> <p><b>NAME</b></p> <p><b>SYNOPSIS</b></p> <pre>#include &lt;stdio.h&gt;  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, ...); ...</pre> <p><b>DESCRIPTION</b></p> <p>The functions in the <i>printf</i>() family produce output according to a <i>format</i> as described below. The function <i>printf</i>() writes output to <i>stdout</i>, the standard output stream; <i>fprintf</i>() writes output to the given output stream; <i>sprintf</i>() and <i>snprintf</i>() write to the character string <i>str</i>.</p> <p>The function <i>snprintf</i>() writes at most <i>size</i> bytes (including the trailing null byte ('\0')) to <i>str</i>.</p> <p>These functions write the output under the control of a <i>format</i> string that specifies how subsequent arguments (or arguments accessed via the variable-length argument facilities of <i>stdarg</i>(3)) are converted for output.</p> <p><b>Return value</b></p> <p>Upon successful return, these functions return the number of characters printed (not including the trailing '\0' used to end output to strings).</p> <p>The functions <i>sprintf</i>() and <i>vsnprintf</i>() do not write more than <i>size</i> 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 <i>size</i> or more means that the output was truncated.</p> <p>If an output error is encountered, a negative value is returned.</p> <p><b>Format of the format string</b></p> <p>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 <i>conversion specifier</i>. In between there may be (in this order) zero or more <i>flags</i>, an optional minimum <i>field width</i>, an optional <i>precision</i> and an optional <i>length modifier</i>.</p> <p><b>The conversion specifier</b></p> <p>A character that specifies the type of conversion to be applied. An example for a conversion specifier is:</p> <p><b>o, u, x, X</b> The <i>unsigned int</i> argument is converted to unsigned octal (o), unsigned decimal (u), or unsigned hexadecimal (x and X) notation.</p> <p><b>s</b> The <i>const char*</i> 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.</p> <p><b>SEE ALSO</b></p> <p><i>printf</i>(1), <i>asprintf</i>(3), <i>dprintf</i>(3), <i>scanf</i>(3), <i>setlocale</i>(3), <i>wctomb</i>(3), <i>wprintf</i>(3), <i>locale</i>(5)</p>
<p>GSP-Klausur Manual-Auszug</p> <p>2020-08-11</p> <p>1</p>	<p>GSP-Klausur Manual-Auszug</p> <p>2020-08-11</p> <p>1</p>

<p>pthread_create(pthread_exit(3))</p> <p>pthread_create – create a new thread / pthread_exit – terminate the calling thread</p> <p><b>NAME</b></p> <p>pthread_create / pthread_exit – terminate the calling thread</p> <p><b>SYNOPSIS</b></p> <pre>#include &lt;pthread.h&gt;  int pthread_create(pthread_t * thread, pthread_attr_t * attr, void * (*start_routine)(void *), void * arg);  void pthread_exit(void *retval);</pre> <p><b>DESCRIPTION</b></p> <p><b>pthread_create</b> creates a new thread of control that executes concurrently with the calling thread. The new thread applies the function <i>start_routine</i> passing it <i>arg</i> as first argument. The new thread terminates either explicitly, by calling <b>pthread_exit(3)</b>, or implicitly, by returning from the <i>start_routine</i> function. The latter case is equivalent to calling <b>pthread_exit(3)</b> with the result returned by <i>start_routine</i> as exit code.</p> <p>The <i>attr</i> argument specifies thread attributes to be applied to the new thread. See <b>pthread_attr_init(3)</b> for a complete list of thread attributes. The <i>attr</i> argument can also be <b>NULL</b>, in which case default attributes are used: the created thread is joinable (not detached) and has default (non real-time) scheduling policy.</p> <p><b>pthread_exit</b> terminates the execution of the calling thread. All cleanup handlers that have been set for the calling thread with <b>pthread_cleanup_push(3)</b> 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-<b>NULL</b> values associated with them in the calling thread (see <b>pthread_key_create(3)</b>). Finally, execution of the calling thread is stopped.</p> <p>The <i>retval</i> argument is the return value of the thread. It can be consulted from another thread using <b>pthread_join(3)</b>.</p> <p><b>RETURN VALUE</b></p> <p>On success, the identifier of the newly created thread is stored in the location pointed by the <i>thread</i> argument, and a 0 is returned. On error, a non-zero error code is returned.</p> <p>The <b>pthread_exit</b> function never returns.</p> <p><b>ERRORS</b></p> <p><b>EAGAIN</b> not enough system resources to create a process for the new thread.</p> <p><b>EAGAIN</b> more than <b>PTHREAD_THREADS_MAX</b> threads are already active.</p> <p><b>AUTHOR</b></p> <p>Xavier Leroy &lt;Xavier.Leroy@inria.fr&gt;</p> <p><b>SEE ALSO</b></p> <p><b>pthread_join(3)</b>, <b>pthread_detach(3)</b>, <b>pthread_attr_init(3)</b>.</p>	<p>pthread_join(3)</p> <p>pthread_join – join with a terminated thread</p> <p><b>NAME</b></p> <p>pthread_join – join with a terminated thread</p> <p><b>SYNOPSIS</b></p> <pre>#include &lt;pthread.h&gt;  int pthread_join(pthread_t thread, void **retval);</pre> <p>Compile and link with <i>-pthread</i>.</p> <p><b>DESCRIPTION</b></p> <p>The <b>pthread_join(0)</b> function waits for the thread specified by <i>thread</i> to terminate. If that thread has already terminated, then <b>pthread_join(0)</b> returns immediately. The thread specified by <i>thread</i> must be joinable.</p> <p>If <i>retval</i> is not <b>NULL</b>, then <b>pthread_join(0)</b> copies the exit status of the target thread (i.e., the value that the target thread supplied to <b>pthread_exit(3)</b>) into the location pointed to by <i>retval</i>. If the target thread was canceled, then <b>PTHREAD_CANCELED</b> is placed in the location pointed to by <i>retval</i>.</p> <p>If multiple threads simultaneously try to join with the same thread, the results are undefined. If the thread calling <b>pthread_join(0)</b> is canceled, then the target thread will remain joinable (i.e., it will not be detached).</p> <p><b>RETURN VALUE</b></p> <p>On success, <b>pthread_join(0)</b> returns 0; on error, it returns an error number.</p> <p><b>ERRORS</b></p> <p><b>EDEADLK</b> A deadlock was detected (e.g., two threads tried to join with each other); or <i>thread</i> specifies the calling thread.</p> <p><b>EINVAL</b> <i>thread</i> is not a joinable thread.</p> <p><b>EINVAL</b> Another thread is already waiting to join with this thread.</p> <p><b>ESRCH</b> No thread with the ID <i>thread</i> could be found.</p> <p><b>NOTES</b></p> <p>After a successful call to <b>pthread_join(0)</b>, the caller is guaranteed that the target thread has terminated. The caller may then choose to do any clean-up that is required after termination of the thread (e.g., freeing memory or other resources that were allocated to the target thread).</p> <p>Joining with a thread that has previously been joined results in undefined behavior.</p> <p>Failure to join with a thread that is joinable (i.e., one that is not detached), produces a "zombie thread". Avoid doing this, since each zombie thread consumes some system resources, and when enough zombie threads have accumulated, it will no longer be possible to create new threads (or processes).</p> <p>There is no pthreads analog of <i>waitpid(-1, &amp;status, 0)</i>, that is, "join with any terminated thread". If you believe you need this functionality, you probably need to rethink your application design.</p> <p>All of the threads in a process are peers: any thread can join with any other thread in the process.</p> <p><b>EXAMPLE</b></p> <p>See <b>pthread_create(3)</b>.</p> <p><b>SEE ALSO</b></p> <p><b>pthread_cancel(3)</b>, <b>pthread_create(3)</b>, <b>pthread_detach(3)</b>, <b>pthread_exit(3)</b>, <b>pthread_t(7)</b></p>	<p>pthread_create(pthread_exit(3))</p> <p>pthread_create – create a new thread / pthread_exit – terminate the calling thread</p> <p><b>NAME</b></p> <p>pthread_create / pthread_exit – terminate the calling thread</p> <p><b>SYNOPSIS</b></p> <pre>#include &lt;pthread.h&gt;  int pthread_create(pthread_t * thread, pthread_attr_t * attr, void * (*start_routine)(void *), void * arg);  void pthread_exit(void *retval);</pre> <p><b>DESCRIPTION</b></p> <p><b>pthread_create</b> creates a new thread of control that executes concurrently with the calling thread. The new thread applies the function <i>start_routine</i> passing it <i>arg</i> as first argument. The new thread terminates either explicitly, by calling <b>pthread_exit(3)</b>, or implicitly, by returning from the <i>start_routine</i> function. The latter case is equivalent to calling <b>pthread_exit(3)</b> with the result returned by <i>start_routine</i> as exit code.</p> <p>The <i>attr</i> argument specifies thread attributes to be applied to the new thread. See <b>pthread_attr_init(3)</b> for a complete list of thread attributes. The <i>attr</i> argument can also be <b>NULL</b>, in which case default attributes are used: the created thread is joinable (not detached) and has default (non real-time) scheduling policy.</p> <p><b>pthread_exit</b> terminates the execution of the calling thread. All cleanup handlers that have been set for the calling thread with <b>pthread_cleanup_push(3)</b> 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-<b>NULL</b> values associated with them in the calling thread (see <b>pthread_key_create(3)</b>). Finally, execution of the calling thread is stopped.</p> <p>The <i>retval</i> argument is the return value of the thread. It can be consulted from another thread using <b>pthread_join(3)</b>.</p> <p><b>RETURN VALUE</b></p> <p>On success, the identifier of the newly created thread is stored in the location pointed by the <i>thread</i> argument, and a 0 is returned. On error, a non-zero error code is returned.</p> <p>The <b>pthread_exit</b> function never returns.</p> <p><b>ERRORS</b></p> <p><b>EAGAIN</b> not enough system resources to create a process for the new thread.</p> <p><b>EAGAIN</b> more than <b>PTHREAD_THREADS_MAX</b> threads are already active.</p> <p><b>AUTHOR</b></p> <p>Xavier Leroy &lt;Xavier.Leroy@inria.fr&gt;</p> <p><b>SEE ALSO</b></p> <p><b>pthread_join(3)</b>, <b>pthread_detach(3)</b>, <b>pthread_attr_init(3)</b>.</p>	<p>pthread_join(3)</p> <p>pthread_join – join with a terminated thread</p> <p><b>NAME</b></p> <p>pthread_join – join with a terminated thread</p> <p><b>SYNOPSIS</b></p> <pre>#include &lt;pthread.h&gt;  int pthread_join(pthread_t thread, void **retval);</pre> <p>Compile and link with <i>-pthread</i>.</p> <p><b>DESCRIPTION</b></p> <p>The <b>pthread_join(0)</b> function waits for the thread specified by <i>thread</i> to terminate. If that thread has already terminated, then <b>pthread_join(0)</b> returns immediately. The thread specified by <i>thread</i> must be joinable.</p> <p>If <i>retval</i> is not <b>NULL</b>, then <b>pthread_join(0)</b> copies the exit status of the target thread (i.e., the value that the target thread supplied to <b>pthread_exit(3)</b>) into the location pointed to by <i>retval</i>. If the target thread was canceled, then <b>PTHREAD_CANCELED</b> is placed in the location pointed to by <i>retval</i>.</p> <p>If multiple threads simultaneously try to join with the same thread, the results are undefined. If the thread calling <b>pthread_join(0)</b> is canceled, then the target thread will remain joinable (i.e., it will not be detached).</p> <p><b>RETURN VALUE</b></p> <p>On success, <b>pthread_join(0)</b> returns 0; on error, it returns an error number.</p> <p><b>ERRORS</b></p> <p><b>EDEADLK</b> A deadlock was detected (e.g., two threads tried to join with each other); or <i>thread</i> specifies the calling thread.</p> <p><b>EINVAL</b> <i>thread</i> is not a joinable thread.</p> <p><b>EINVAL</b> Another thread is already waiting to join with this thread.</p> <p><b>ESRCH</b> No thread with the ID <i>thread</i> could be found.</p> <p><b>NOTES</b></p> <p>After a successful call to <b>pthread_join(0)</b>, the caller is guaranteed that the target thread has terminated. The caller may then choose to do any clean-up that is required after termination of the thread (e.g., freeing memory or other resources that were allocated to the target thread).</p> <p>Joining with a thread that has previously been joined results in undefined behavior.</p> <p>Failure to join with a thread that is joinable (i.e., one that is not detached), produces a "zombie thread". Avoid doing this, since each zombie thread consumes some system resources, and when enough zombie threads have accumulated, it will no longer be possible to create new threads (or processes).</p> <p>There is no pthreads analog of <i>waitpid(-1, &amp;status, 0)</i>, that is, "join with any terminated thread". If you believe you need this functionality, you probably need to rethink your application design.</p> <p>All of the threads in a process are peers: any thread can join with any other thread in the process.</p> <p><b>EXAMPLE</b></p> <p>See <b>pthread_create(3)</b>.</p> <p><b>SEE ALSO</b></p> <p><b>pthread_cancel(3)</b>, <b>pthread_create(3)</b>, <b>pthread_detach(3)</b>, <b>pthread_exit(3)</b>, <b>pthread_t(7)</b></p>
<p>GSP-Klausur Manual-Auszug</p> <p>2020-08-11</p>	<p>GSP-Klausur Manual-Auszug</p> <p>2020-08-11</p>	<p>GSP-Klausur Manual-Auszug</p> <p>2020-08-11</p>	<p>GSP-Klausur Manual-Auszug</p> <p>2020-08-11</p>

stat(2)

stat(2)

stat(2)

stat(2)

**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)**

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.