

accept(2)	<p>accept – accept a connection on a socket</p> <p>NAME</p> <p>accept – accept a connection on a socket</p> <p>SYNOPSIS</p> <pre>#include <sys/types.h> #include <sys/socket.h> int accept(int s, struct sockaddr *addr, int *addrlen);</pre> <p>DESCRIPTION</p> <p>The argument <i>s</i> is a socket that has been created with <code>socket(3N)</code> and bound to an address with <code>bind(3N)</code>, and that is listening for connections; after a call to <code>listen(3N)</code>. The <code>accept()</code> function extracts the first connection on the queue of pending connections, creates a new socket with the properties of <i>s</i>, and allocates a new file descriptor, <i>ns</i>, for the socket. If no pending connections are present on the queue and the socket is not marked as non-blocking, <code>accept()</code> blocks the caller until a connection is present. If the socket is marked as non-blocking and no pending connections are present on the queue, <code>accept()</code> returns an error as described below. The <code>accept()</code> function uses the <code>netconfig(4)</code> file to determine the STREAMS device file name associated with <i>s</i>. This is the device on which the connect indication will be accepted. The accepted socket, <i>ns</i>, is used to read and write data to and from the socket that connected to <i>ns</i>; it is not used to accept more connections. The original socket (<i>s</i>) remains open for accepting further connections.</p> <p>The argument <i>addr</i> is a result parameter that is filled in with the address of the connecting entity as it is known to the communications layer. The exact format of the <i>addr</i>-parameter is determined by the domain in which the communication occurs.</p> <p>The argument <i>addrLen</i> is a value-result parameter. Initially, it contains the amount of space pointed to by <i>addr</i>; on return it contains the length in bytes of the address returned.</p> <p>The <code>accept()</code> function is used with connection-based socket types, currently with <code>SOCK_STREAM</code>.</p> <p>It is possible to <code>select(3C)</code> or <code>poll(2)</code> a socket for the purpose of an <code>accept()</code> by selecting or polling it for a read. However, this will only indicate when a connect indication is pending, it is still necessary to call <code>accept()</code>.</p> <p>RETURN VALUE</p> <p>On success, these system calls return a nonnegative integer that is a file descriptor for the accepted socket. On error, <code>-1</code> is returned, and <i>errno</i> is set appropriately.</p> <p>ERRORS</p> <p><code>accept()</code> will fail if:</p> <ul style="list-style-type: none"> EBADF The descriptor is invalid. EINTR The accept attempt was interrupted by the delivery of a signal. EMFILE The per-process descriptor table is full. ENODEV The protocol family and type corresponding to <i>s</i> could not be found in the <code>netconfig</code> file. ENOMEM There was insufficient user memory available to complete the operation. EPROTO A protocol error has occurred; for example, the STREAMS protocol stack has not been initialized or the connection has already been released. EWOLDBLOCK The socket is marked as non-blocking and no connections are present to be accepted. <p>SEE ALSO</p> <p><code>poll(2)</code>, <code>bind(3N)</code>, <code>connect(3N)</code>, <code>select(3C)</code>, <code>socket(3N)</code>, <code>netconfig(4)</code>, <code>attributes(5)</code>, <code>socket(5)</code></p>	accept(2)	SP-Klausur Manual-Auszug	2021-02-24	1
bind(2)	<p>bind – bind a name to a socket</p> <p>NAME</p> <p>bind – bind a name to a socket</p> <p>SYNOPSIS</p> <pre>#include <sys/types.h> #include <sys/socket.h> int bind(int s, const struct sockaddr *name, int namelen);</pre> <p>DESCRIPTION</p> <p><code>bind()</code> assigns a name to an unnamed socket <i>s</i>. When a socket is created with <code>socket(3N)</code>, it exists in a name space (address family) but has no name assigned. <code>bind()</code> requests that the name pointed to by <i>name</i> be assigned to the socket.</p> <p>RETURN VALUE</p> <p>On success, zero is returned. On error, <code>-1</code> is returned, and <i>errno</i> is set appropriately.</p> <p>ERRORS</p> <p>The <code>bind()</code> call will fail if:</p> <ul style="list-style-type: none"> EACCES The requested address is protected and the current user has inadequate permission to access it. EADDRINUSE The specified address is already in use. EADDRNOTAVAIL The specified address is not available on the local machine. EBADF <i>s</i> is not a valid descriptor. EINVAL <i>namelen</i> is not the size of a valid address for the specified address family. EINVAL The socket is already bound to an address. ENOSR There were insufficient STREAMS resources for the operation to complete. ENOTSOCK <i>s</i> is a descriptor for a file, not a socket. <p>The following errors are specific to binding names in the UNIX domain:</p> <ul style="list-style-type: none"> EACCES Search permission is denied for a component of the path prefix of the pathname in <i>name</i>. EIO An I/O error occurred while making the directory entry or allocating the inode. EISDIR A null pathname was specified. ELOOP Too many symbolic links were encountered in translating the pathname in <i>name</i>. ENOENT A component of the path prefix of the pathname in <i>name</i> does not exist. ENOTDIR A component of the path prefix of the pathname in <i>name</i> is not a directory. EROFS The inode would reside on a read-only file system. <p>SEE ALSO</p> <p><code>unlink(2)</code>, <code>socket(3N)</code>, <code>attributes(5)</code>, <code>socket(5)</code></p> <p>NOTES</p> <p>Binding a name in the UNIX domain creates a socket in the file system that must be deleted by the caller when it is no longer needed (using <code>unlink(2)</code>).</p> <p>The rules used in name binding vary between communication domains.</p>	bind(2)	SP-Klausur Manual-Auszug	2021-02-24	1

close(2)	<p>close – close a file descriptor</p> <p>NAME</p> <p>close – close a file descriptor</p> <p>SYNOPSIS</p> <pre>#include <unistd.h> int close(int <i>fd</i>);</pre> <p>DESCRIPTION</p> <p><code>close()</code> closes a file descriptor, so that it no longer refers to any file and may be reused. Any record locks (see <code>fcntl(2)</code>) held on the file it was associated with, and owned by the process, are removed (regardless of the file descriptor that was used to obtain the lock).</p> <p>If <i>fd</i> is the last file descriptor referring to the underlying open file description (see <code>open(2)</code>), the resources associated with the open file description are freed; if the file descriptor was the last reference to a file which has been removed using <code>unlink(2)</code>, the file is deleted.</p> <p>RETURN VALUE</p> <p><code>close()</code> returns zero on success. On error, <code>-1</code> is returned, and <code>errno</code> is set appropriately.</p> <p>ERRORS</p> <p>EBADF <i>fd</i> isn't a valid open file descriptor.</p> <p>EINTR The <code>close()</code> call was interrupted by a signal; see <code>signal(7)</code>.</p> <p>EIO An I/O error occurred.</p> <p>ENOSPC, EDQUOT On NFS, these errors are not normally reported against the first write which exceeds the available storage space, but instead against a subsequent <code>write(2)</code>, <code>fsync(2)</code>, or <code>close(2)</code>.</p>	close(2)
dup(2)	<p>dup, dup2 – duplicate a file descriptor</p> <p>NAME</p> <p>dup, dup2 – duplicate a file descriptor</p> <p>SYNOPSIS</p> <pre>#include <unistd.h> int dup(int <i>oldfd</i>); int dup2(int <i>oldfd</i>, int <i>newfd</i>);</pre> <p>DESCRIPTION</p> <p><code>dup()</code> and <code>dup2()</code> create a copy of the file descriptor <i>oldfd</i>.</p> <p><code>dup()</code> uses the lowest-numbered unused descriptor for the new descriptor.</p> <p><code>dup2()</code> makes <i>newfd</i> be the copy of <i>oldfd</i>, closing <i>newfd</i> first if necessary, but note the following:</p> <ul style="list-style-type: none"> * If <i>oldfd</i> is not a valid file descriptor, then the call fails, and <i>newfd</i> is not closed. * If <i>oldfd</i> is a valid file descriptor, and <i>newfd</i> has the same value as <i>oldfd</i>, then <code>dup2()</code> does nothing, and returns <i>newfd</i>. <p>After a successful return from <code>dup()</code> or <code>dup2()</code>, the old and new file descriptors may be used interchangeably. They refer to the same open file description (see <code>open(2)</code>) and thus share file offset and file status flags; for example, if the file offset is modified by using <code>lseek(2)</code> on one of the descriptors, the offset is also changed for the other.</p> <p>The two descriptors do not share file descriptor flags (the close-on-exec flag). The close-on-exec flag (<code>FD_CLOEXEC</code>; see <code>fcntl(2)</code>) for the duplicate descriptor is off.</p> <p>RETURN VALUE</p> <p><code>dup()</code> and <code>dup2()</code> return the new descriptor, or <code>-1</code> if an error occurred (in which case, <code>errno</code> is set appropriately).</p> <p>ERRORS</p> <p>EBADF <i>oldfd</i> isn't an open file descriptor, or <i>newfd</i> is out of the allowed range for file descriptors.</p> <p>EBUSY (Linux only) This may be returned by <code>dup2()</code> during a race condition with <code>open(2)</code> and <code>dup()</code>.</p> <p>EINTR The <code>dup2()</code> call was interrupted by a signal; see <code>signal(7)</code>.</p> <p>EMFILE The process already has the maximum number of file descriptors open and tried to open a new one.</p> <p>SEE ALSO <code>close(2)</code>, <code>fcntl(2)</code>, <code>open(2)</code></p>	dup(2)
SP-Klausur Manual-Auszug	2021-02-24	1
SP-Klausur Manual-Auszug	2021-02-24	1

NAME

clearerr, feof, ferorr, fileno – check and reset stream status

SYNOPSIS

```
#include <stdio.h>
```

```
void clearerr(FILE *stream);
int feof(FILE *stream);
int ferorr(FILE *stream);
int fileno(FILE *stream);
```

DESCRIPTION

The function **clearerr()** clears the end-of-file and error indicators for the stream pointed to by *stream*.

The function **feof()** tests the end-of-file indicator for the stream pointed to by *stream*, returning non-zero if it is set. The end-of-file indicator can only be cleared by the function **clearerr()**.

The function **ferorr()** tests the error indicator for the stream pointed to by *stream*, returning non-zero if it is set. The error indicator can only be reset by the **clearerr()** function.

The function **fileno()** examines the argument *stream* and returns its integer descriptor.

For non-locking counterparts, see **unlocked_stdio(3)**.

ERRORS

These functions should not fail and do not set the external variable *errno*. (However, in case **fileno()** detects that its argument is not a valid stream, it must return -1 and set *errno* to **EBADF**.)

CONFORMING TO

The functions **clearerr()**, **feof()**, and **ferorr()** conform to C89 and C99.

SEE ALSO

open(2), **fdopen(3)**, **stdio(3)**, **unlocked_stdio(3)**

NAME

fopen, fdopen, fileno – stream open functions

SYNOPSIS

```
#include <stdio.h>
```

```
FILE *fopen(const char *path, const char *mode);
FILE *fdopen(int fd, const char *mode);
int fileno(FILE *stream);
int fclose(FILE *stream);
```

DESCRIPTION

The **fopen** function opens the file whose name is the string pointed to by *path* and associates a stream with it.

The argument *mode* points to a string beginning with one of the following sequences (Additional characters may follow these sequences.):

r Open text file for reading. The stream is positioned at the beginning of the file.

r+ Open for reading and writing. The stream is positioned at the beginning of the file.

w Truncate file to zero length or create text file for writing. The stream is positioned at the beginning of the file.

w+ Open for reading and writing. The file is created if it does not exist, otherwise it is truncated. The stream is positioned at the beginning of the file.

a Open for appending (writing at end of file). The file is created if it does not exist. The stream is positioned at the end of the file.

a+ Open for reading and appending (writing at end of file). The file is created if it does not exist. The stream is positioned at the end of the file.

The **fdopen** function associates a stream with the existing file descriptor, *fd*. The *mode* of the stream (one of the values "r", "r+", "w", "w+", "a", "a+") must be compatible with the mode of the file descriptor. The file position indicator of the new stream is set to that belonging to *fd*, and the error and end-of-file indicators are cleared. Modes "w" or "w+" do not cause truncation of the file. The file descriptor is not dup'ed, and will be closed when the stream created by **fdopen** is closed. The result of applying **fdopen** to a shared memory object is undefined.

The function **fileno()** examines the argument *stream* and returns its integer descriptor.

The **fclose()** function flushes the stream pointed to by *stream* (writing any buffered output data using **flush(3)**) and closes the underlying file descriptor.

RETURN VALUE

Upon successful completion **fopen**, **fdopen** and **freopen** return a **FILE** pointer. Otherwise, **NULL** is returned and the global variable *errno* is set to indicate the error. Upon successful completion of **fclose**, 0 is returned. Otherwise, **EOF** is returned and *errno* is set to indicate the error.

ERRORS

EINVAL

The *mode* provided to **fopen**, **fdopen**, or **freopen** was invalid.

EBADF

The file descriptor underlying *stream* passed to **fclose** is not valid.

The **fopen**, **fdopen** and **freopen** functions may also fail and set *errno* for any of the errors specified for the routine **malloc(3)**.

The **fopen** function may also fail and set *errno* for any of the errors specified for the routine **open(2)**.

The **fdopen** function may also fail and set *errno* for any of the errors specified for the routine **fcntl(2)**.

<p>getc/fgets/putc/fputs(3)</p> <p>NAME fgetc, fgets, getc, getchar, fputc, fputs, putc, putchar – input and output of characters and strings</p> <p>SYNOPSIS #include <stdio.h></p> <pre>int fgetc(FILE *stream); char *fgets(char *s, int size, FILE *stream); int getc(FILE *stream); int getchar(void); int fputc(int c, FILE *stream); int fputs(const char *s, FILE *stream); int puts(int c, FILE *stream); int putchar(int c);</pre> <p>DESCRIPTION fgetc() reads the next character from <i>stream</i> and returns it as an <i>unsigned char</i> cast to an <i>int</i>, or EOF on end of file or error. getc() is equivalent to fgetc() except that it may be implemented as a macro which evaluates <i>stream</i> more than once. getchar() is equivalent to getc(stdin). fgets() reads in at most one less than <i>size</i> characters from <i>stream</i> and stores them into the buffer pointed to by <i>s</i>. Reading stops after an EOF or a newline. If a newline is read, it is stored into the buffer. A '\0' is stored after the last character in the buffer. fputc() writes the character <i>c</i>, cast to an <i>unsigned char</i>, to <i>stream</i>. fputs() writes the string <i>s</i> to <i>stream</i>, without its terminating null byte ('\0'). putc() is equivalent to fputc() except that it may be implemented as a macro which evaluates <i>stream</i> more than once. putchar(c); is equivalent to putc(c, stdout). Calls to the functions described here can be mixed with each other and with calls to other output functions from the <i>stdio</i> library for the same output stream.</p> <p>RETURN VALUE fgetc(), getc() and getchar() return the character read as an <i>unsigned char</i> cast to an <i>int</i> or EOF on end of file or error. fgets() returns <i>s</i> on success, and NULL on error or when end of file occurs while no characters have been read. fputc(), putc() and putchar() return the character written as an <i>unsigned char</i> cast to an <i>int</i> or EOF on error. fputs() returns a nonnegative number on success, or EOF on error.</p> <p>SEE ALSO read(2), write(2), ferrror(3), fgetcws(3), fgets(3), fopen(3), fread(3), fseek(3), fgetc(3), getwchar(3), scanf(3), ungetc(3), write(2), ferrror(3), fopen(3), fopenc(3), fputcws(3), fputs(3), fseek(3), fwrite(3), gets(3), putchar(3), scanf(3), ungetc(3), unlocked_stdio(3)</p>	<p>ip(v6)/socket(7)</p> <p>NAME ip(v6, AF_INET6 – Linux IPv6 protocol implementation</p> <p>SYNOPSIS #include <sys/socket.h> #include <netinet/in.h></p> <pre>tcp6_socket = socket(AF_INET6, SOCK_STREAM, 0); raw6_socket = socket(AF_INET6, SOCK_RAW, protocol); udp6_socket = socket(AF_INET6, SOCK_DGRAM, protocol);</pre> <p>DESCRIPTION Linux 2.2 optionally implements the Internet Protocol, version 6. This man page contains a description of the IPv6 basic API as implemented by the Linux kernel and glibc 2.1. The interface is based on the BSD sockets interface; see socket(7). The IPv6 API aims to be mostly compatible with the ip(7) v4 API. Only differences are described in this man page. To bind an AF_INET6 socket to any process the local address should be copied from the <i>in6addr_*</i> variable which has <i>in6_addr</i> type. In static initializations IN6ADDR_ANY_INIT may also be used, which expands to a constant expression. Both of them are in network order. IPv4 connections can be handled with the v6 API by using the v4-mapped-on-v6 address type; thus a program only needs only to support this API type to support both protocols. This is handled transparently by the address handling functions in libc. IPv4 and IPv6 share the local port space. When you get an IPv4 connection or packet to a IPv6 socket its source address will be mapped to v6 and it will be mapped to v6.</p> <p>Address Format</p> <pre>struct sockaddrr_in6 { uint16_t sin6_family; /* AF_INET6 */ uint16_t sin6_port; /* port number */ uint32_t sin6_flowinfo; /* IPv6 flow information */ struct in6_addr sin6_addr; /* IPv6 address */ uint32_t sin6_scope_id; /* Scope ID (new in 2.4) */ }; struct in6_addr { unsigned char s6_addr[16]; /* IPv6 address */ };</pre> <p><i>sin6_family</i> is always set to AF_INET6; <i>sin6_port</i> is the protocol port (see <i>sin_port</i> in ip(7)); <i>sin6_flowinfo</i> is the IPv6 flow identifier; <i>sin6_addr</i> is the 128-bit IPv6 address. <i>sin6_scope_id</i> is an ID of depending of the scope of the address. It is new in Linux 2.4. Linux only supports it for link scope addresses, in that case <i>sin6_scope_id</i> contains the interface index (see netdevice(7))</p> <p>RETURN VALUES –1 is returned if an error occurs. Otherwise the return value is a descriptor referencing the socket.</p> <p>NOTES The <i>sockaddrr_in6</i> structure is bigger than the generic <i>sockaddr</i>. Programs that assume that all address types can be stored safely in a <i>struct sockaddr</i> need to be changed to use <i>struct sockaddrr_storage</i> for that instead.</p> <p>SEE ALSO cmsg(3), ip(7)</p>	<p>ip(v6)/socket(7)</p>
<p>SP-Klausur Manual-Auszug</p> <p>2021-02-24</p>	<p>SP-Klausur Manual-Auszug</p> <p>2021-02-24</p>	<p>SP-Klausur Manual-Auszug</p> <p>2021-02-24</p>

listen(2)	listen – listen for connections on a socket	opendir/readdir(3)	opendir – open a directory / readdir – read a directory
NAME	listen – listen for connections on a socket	opendir/readdir(3)	opendir – open a directory / readdir – read a directory
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/socket.h></pre>	opendir/readdir(3)	<pre>#include <sys/types.h> #include <dirent.h></pre>
DESCRIPTION	<p><code>int listen(int sockfd, int backlog);</code></p> <p><code>listen()</code> marks the socket referred to by <code>sockfd</code> as a passive socket, that is, as a socket that will be used to accept incoming connection requests using <code>accept(2)</code>.</p> <p>The <code>sockfd</code> argument is a file descriptor that refers to a socket of type SOCK_STREAM or SOCK_SEQPACKET.</p> <p>The <code>backlog</code> argument defines the maximum length to which the queue of pending connections for <code>sockfd</code> may grow. If a connection request arrives when the queue is full, the client may receive an error with an indication of ECONNREFUSED or, if the underlying protocol supports retransmission, the request may be ignored so that a later reattempt at connection succeeds.</p>	opendir/readdir(3)	<pre>DIR *opendir(const char *name); int closedir(DIR *dir); struct dirent *readdir(DIR *dir);</pre> <p>The <code>opendir()</code> function opens a directory stream corresponding to the directory <code>name</code>, and returns a pointer to the directory stream. The stream is positioned at the first entry in the directory.</p> <p>RETURN VALUE</p> <p>The <code>opendir()</code> function returns a pointer to the directory stream. On error, <code>NULL</code> is returned, and <code>errno</code> is set appropriately.</p> <p>DESCRIPTION closedir</p> <p>The <code>closedir()</code> function closes the directory stream associated with <code>dirp</code>. A successful call to <code>closedir()</code> also closes the underlying file descriptor associated with <code>dirp</code>. The directory stream descriptor <code>dirp</code> is not available after this call.</p> <p>RETURN VALUE</p> <p>The <code>closedir()</code> function returns 0 on success. On error, <code>-1</code> is returned, and <code>errno</code> is set appropriately.</p> <p>DESCRIPTION readdir</p> <p>The <code>readdir()</code> function returns a pointer to a dirent structure representing the next directory entry in the directory stream pointed to by <code>dir</code>. It returns <code>NULL</code> 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 <code>dir</code> are created by distinct calls to <code>opendir()</code>.</p> <p>The data returned by <code>readdir()</code> is overwritten by subsequent calls to <code>readdir()</code> for the same directory stream.</p> <p>The <code>dirent</code> structure is defined as follows:</p> <pre>struct dirent { long d_ino; /* inode number */ char d_name[256]; /* filename */ };</pre> <p>RETURN VALUE</p> <p>On success, <code>readdir()</code> returns a pointer to a <code>dirent</code> structure. (This structure may be statically allocated; do not attempt to <code>free(3)</code> it.)</p> <p>If the end of the directory stream is reached, <code>NULL</code> is returned and <code>errno</code> is not changed. If an error occurs, <code>NULL</code> is returned and <code>errno</code> is set appropriately. To distinguish end of stream and from an error, set <code>errno</code> to zero before calling <code>readdir()</code> and then check the value of <code>errno</code> if <code>NULL</code> is returned.</p> <p>ERRORS</p> <p>EACCES Permission denied.</p> <p>ENOENT Directory does not exist, or <code>name</code> is an empty string.</p> <p>ENOTDIR <code>name</code> is not a directory.</p>
NOTES	<p>To accept connections, the following steps are performed:</p> <ol style="list-style-type: none"> 1. A socket is created with <code>socket(2)</code>. 2. The socket is bound to a local address using <code>bind(2)</code>, so that other sockets may be <code>connect(2)</code>ed to it. 3. A willingness to accept incoming connections and a queue limit for incoming connections are specified with <code>listen()</code>. 4. Connections are accepted with <code>accept(2)</code>. <p>If the <code>backlog</code> argument is greater than the value in <code>/proc/sys/net/core/somaxconn</code>, then it is silently truncated to that value; the default value in this file is 128.</p> <p>EXAMPLE</p> <p>See <code>bind(2)</code>.</p> <p>SEE ALSO</p> <p><code>accept(2)</code>, <code>bind(2)</code>, <code>connect(2)</code>, <code>socket(2)</code>, <code>socket(7)</code></p>	opendir/readdir(3)	opendir – open a directory / readdir – read a directory
SP-Klausur Manual-Auszug	2021-02-24	SP-Klausur Manual-Auszug	2021-02-24

pthread_cond(3)	pthread_cond(3)	pthread_cond(3)
<p>NAME</p> <p>pthread_cond_init, pthread_cond_destroy, pthread_cond_signal, pthread_cond_wait, pthread_cond_timedwait – operations on conditions</p> <p>SYNOPSIS</p> <pre>#include <pthread.h> pthread_cond_t cond = PTHREAD_COND_INITIALIZER; int pthread_cond_init(pthread_cond_t *cond, pthread_condattr_t *cond_attr); int pthread_cond_signal(pthread_cond_t *cond); int pthread_cond_broadcast(pthread_cond_t *cond); int pthread_cond_wait(pthread_cond_t *cond, pthread_mutex_t *mutex); int pthread_cond_timedwait(pthread_cond_t *cond, pthread_mutex_t *mutex, const struct timespec *abstime); int pthread_cond_destroy(pthread_cond_t *cond);</pre> <p>DESCRIPTION</p> <p>A condition (short for “condition variable”) is a synchronization device that allows threads to suspend execution and relinquish the processors until some predicate on shared data is satisfied. The basic operations on conditions are: signal the condition (when the predicate becomes true), and wait for the condition, suspending the thread execution until another thread signals the condition.</p> <p>A condition variable must always be associated with a mutex, to avoid the race condition where a thread prepares to wait on a condition variable and another thread signals the condition just before the first thread actually waits on it.</p> <p>pthread_cond_init initializes the condition variable <i>cond</i>, using the condition attributes specified in <i>cond_attr</i>, or default attributes if <i>cond_attr</i> is NULL. The LinuxThreads implementation supports no attributes for conditions, hence the <i>cond_attr</i> parameter is actually ignored.</p> <p>Variables of type pthread_cond_t can also be initialized statically, using the constant PTHREAD_COND_INITIALIZER.</p> <p>pthread_cond_signal restarts one of the threads that are waiting on the condition variable <i>cond</i>. If no threads are waiting on <i>cond</i>, nothing happens. If several threads are waiting on <i>cond</i>, exactly one is restarted, but it is not specified which.</p> <p>pthread_cond_broadcast restarts all the threads that are waiting on the condition variable <i>cond</i>. Nothing happens if no threads are waiting on <i>cond</i>.</p> <p>pthread_cond_wait atomically unlocks the <i>mutex</i> (as per pthread_unlock_mutex) and waits for the condition variable <i>cond</i> to be signaled. The thread execution is suspended and does not consume any CPU time until the condition variable is signaled. The <i>mutex</i> must be locked by the calling thread on entrance to pthread_cond_wait. Before returning to the calling thread, pthread_cond_wait re-acquires <i>mutex</i> (as per pthread_lock_mutex).</p> <p>Unlocking the mutex and suspending on the condition variable is done atomically. Thus, if all threads always acquire the mutex before signaling the condition, this guarantees that the condition cannot be</p>	<p>signaled (and thus ignored) between the time a thread locks the mutex and the time it waits on the condition variable.</p> <p>pthread_cond_timedwait atomically unlocks <i>mutex</i> and waits on <i>cond</i>, as pthread_cond_wait does, but it also bounds the duration of the wait. If <i>cond</i> has not been signaled within the amount of time specified by <i>abstime</i>, the mutex <i>mutex</i> is re-acquired and pthread_cond_timedwait returns the error ETIMEOUT. The <i>abstime</i> parameter specifies an absolute time, with the same origin as time(2) and gettimeofday(2): an <i>abstime</i> of 0 corresponds to 00:00:00 GMT, January 1, 1970.</p> <p>pthread_cond_destroy destroys a condition variable, freeing the resources it might hold. No threads must be waiting on the condition variable on entrance to pthread_cond_destroy. In the LinuxThreads implementation, no resources are associated with condition variables, thus pthread_cond_destroy actually does nothing except checking that the condition has no waiting threads.</p> <p>CANCELLATION</p> <p>pthread_cond_wait and pthread_cond_timedwait are cancellation points. If a thread is cancelled while suspended in one of these functions, the thread immediately resumes execution, then locks again the <i>mutex</i> argument to pthread_cond_wait and pthread_cond_timedwait, and finally executes the cancellation. Consequently, cleanup handlers are assured that <i>mutex</i> is locked when they are called.</p> <p>ASYNC-SIGNAL SAFETY</p> <p>The condition functions are not async-signal safe, and should not be called from a signal handler. In particular, calling pthread_cond_signal or pthread_cond_broadcast from a signal handler may deadlock the calling thread.</p> <p>RETURN VALUE</p> <p>All condition variable functions return 0 on success and a non-zero error code on error.</p> <p>ERRORS</p> <p>pthread_cond_init, pthread_cond_signal, pthread_cond_broadcast, and pthread_cond_wait never return an error code.</p> <p>The pthread_cond_timedwait function returns the following error codes on error:</p> <p>ETIMEOUT the condition variable was not signaled until the timeout specified by <i>abstime</i></p> <p>EINTR pthread_cond_timedwait was interrupted by a signal</p> <p>The pthread_cond_destroy function returns the following error code on error:</p> <p>EBUSY some threads are currently waiting on <i>cond</i>.</p> <p>AUTHOR Xavier Leroy <Xavier.Leroy@inria.fr></p> <p>SEE ALSO pthread_condattr_init(3), pthread_mutex_lock(3), pthread_mutex_unlock(3), gettimeofday(2), nanosleep(2).</p>	<p>SP-Klausur Manual-Auszug</p> <p>2021-02-24</p> <p>2</p>
pthread_cond(3)	pthread_cond(3)	pthread_cond(3)
<p>SP-Klausur Manual-Auszug</p> <p>2021-02-24</p> <p>1</p>	<p>SP-Klausur Manual-Auszug</p> <p>2021-02-24</p> <p>2</p>	<p>pthread_cond(3)</p>

pthread_mutex_exit(3) pthread_create(pthread_exit(3) pthread_mutex(3)

NAME
pthread_mutex_init, pthread_mutex_lock, pthread_mutex_trylock, pthread_mutex_unlock, pthread_mutex_destroy – operations on mutexes

SYNOPSIS
#include <pthread.h>
pthread_mutex_t fastmutex = PTHREAD_MUTEX_INITIALIZER;
pthread_mutex_t recursive = PTHREAD_RECURSIVE_MUTEX_INITIALIZER_NP;
pthread_mutex_t errorcheck = PTHREAD_ERRORCHECK_MUTEX_INITIALIZER_NP;
int pthread_mutex_init(pthread_mutex_t *mutex, const pthread_mutexattr_t *mutexattr);
int pthread_mutex_lock(pthread_mutex_t *mutex);
int pthread_mutex_trylock(pthread_mutex_t *mutex);
int pthread_mutex_unlock(pthread_mutex_t *mutex);
int pthread_mutex_destroy(pthread_mutex_t *mutex);

DESCRIPTION
A mutex is a MUTual EXclusion device, and is useful for protecting shared data structures from concurrent modifications, and implementing critical sections and monitors.
A mutex has two possible states: unlocked (not owned by any thread), and locked (owned by one thread). A mutex can never be owned by two different threads simultaneously. A thread attempting to lock a mutex that is already locked by another thread is suspended until the owning thread unlocks the mutex first.
pthread_mutex_init initializes the mutex object pointed to by *mutex* according to the mutex attributes specified in *mutexattr*. If *mutexattr* is **NULL**, default attributes are used instead.
The LinuxThreads implementation supports only one mutex attributes, the *mutex kind*, which is either “fast”, “recursive”, or “error checking”. The kind of a mutex determines whether it can be locked again by a thread that already owns it. The default kind is “fast”. See **pthread_mutexattr_init(3)** for more information on mutex attributes.
Variables of type **pthread_mutex_t** can also be initialized statically, using the constants **PTHREAD_MUTEX_INITIALIZER** (for fast mutexes), **PTHREAD_RECURSIVE_MUTEX_INITIALIZER_NP** (for recursive mutexes), and **PTHREAD_ERRORCHECK_MUTEX_INITIALIZER_NP** (for error checking mutexes).
pthread_mutex_lock locks the given mutex. If the mutex is currently unlocked, it becomes locked and owned by the calling thread, and **pthread_mutex_lock** returns immediately. If the mutex is already locked by another thread, **pthread_mutex_lock** suspends the calling thread until the mutex is unlocked.
If the mutex is already locked by the calling thread, the behavior of **pthread_mutex_lock** depends on the kind of the mutex. If the mutex is of the “fast” kind, the calling thread is suspended until the mutex is unlocked, thus effectively causing the calling thread to deadlock. If the mutex is of the “error checking” kind, **pthread_mutex_lock** returns immediately with the error code **EDEADLK**. If the mutex is of the “recursive” kind, **pthread_mutex_lock** succeeds and returns immediately, recording the number of times the calling thread has locked the mutex. An equal number of **pthread_mutex_unlock** operations must be

pthread_create(pthread_exit(3) pthread_create(pthread_exit(3) pthread_create(pthread_exit(3))

NAME
pthread_create – create a new thread / pthread_exit – terminate the calling thread

SYNOPSIS
#include <pthread.h>
int pthread_create(pthread_t * thread, pthread_attr_t * attr, void * (*start_routine)(void *), void * arg);
void pthread_exit(void *retval);

DESCRIPTION
pthread_create creates a new thread that executes concurrently with the calling thread. The new thread applies the function *start_routine* passing it *arg* as first argument. The new thread terminates either explicitly, by calling **pthread_exit(3)**, or implicitly, by returning from the *start_routine* function. The latter case is equivalent to calling **pthread_exit(3)** with the result returned by *start_routine* as exit code.
The *attr* argument specifies thread attributes to be applied to the new thread. See **pthread_attr_init(3)** for a complete list of thread attributes. The *attr* argument can also be **NULL**, in which case default attributes are used: the created thread is joinable (not detached) and has default (non real-time) scheduling policy.
pthread_exit terminates the execution of the calling thread. All cleanup handlers that have been set for the calling thread with **pthread_cleanup_push(3)** 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-**NULL** values associated with them in the calling thread (see **pthread_key_create(3)**). Finally, execution of the calling thread is stopped.
The *retval* argument is the return value of the thread. It can be consulted from another thread using **pthread_join(3)**.
RETURN VALUE
On success, the identifier of the newly created thread is stored in the location pointed by the *thread* argument, and a 0 is returned. On error, a non-zero error code is returned.
The **pthread_exit** function never returns.
ERRORS
EAGAIN
not enough system resources to create a process for the new thread.
EAGAIN
more than **PTHREAD_THREADS_MAX** threads are already active.
AUTHOR
Xavier Leroy <Xavier.Leroy@inria.fr>
SEE ALSO
pthread_join(3), **pthread_detach(3)**, **pthread_attr_init(3)**.

pthread_mutex(3)	pthread_mutex(3)	sigaction(2)	sigaction(2)
<p>performed before the mutex returns to the unlocked state.</p> <p>pthread_mutex_trylock behaves identically to pthread_mutex_lock, except that it does not block the calling thread if the mutex is already locked by another thread (or by the calling thread in the case of a “fast” mutex). Instead, pthread_mutex_trylock returns immediately with the error code EBUSY.</p> <p>pthread_mutex_unlock unlocks the given mutex. The mutex is assumed to be locked and owned by the calling thread on entrance to pthread_mutex_unlock. If the mutex is of the “fast” kind, pthread_mutex_unlock always returns it to the unlocked state. If it is of the “recursive” kind, it decrements the locking count of the mutex (number of pthread_mutex_lock operations performed on it by the calling thread), and only when this count reaches zero is the mutex actually unlocked.</p> <p>On “error checking” mutexes, pthread_mutex_unlock actually checks at run-time that the mutex is locked on entrance, and that it was locked by the same thread that is now calling pthread_mutex_unlock. If these conditions are not met, an error code is returned and the mutex remains unchanged. “Fast” and “recursive” mutexes perform no such checks, thus allowing a locked mutex to be unlocked by a thread other than its owner. This is non-portable behavior and must not be relied upon.</p> <p>pthread_mutex_destroy destroys a mutex object, freeing the resources it might hold. The mutex must be unlocked on entrance. In the LinuxThreads implementation, no resources are associated with mutex objects, thus pthread_mutex_destroy actually does nothing except checking that the mutex is unlocked.</p> <p>RETURN VALUE pthread_mutex_init always returns 0. The other mutex functions return 0 on success and a non-zero error code on error.</p> <p>ERRORS The pthread_mutex_lock function returns the following error code on error: EINVAL the mutex has not been properly initialized. EDEADLK the mutex is already locked by the calling thread (“error checking” mutexes only). The pthread_mutex_unlock function returns the following error code on error: EINVAL the mutex has not been properly initialized. EPERM the calling thread does not own the mutex (“error checking” mutexes only). The pthread_mutex_destroy function returns the following error code on error: EBUSY the mutex is currently locked.</p> <p>AUTHOR Xavier Leroy <Xavier.Leroy@inria.fr></p> <p>SEE ALSO pthread_mutexattr_init(3), pthread_mutexattr_setkind_np(3), pthread_cancel(3).</p>	<p>NAME sigaction – POSIX signal handling functions.</p> <p>SYNOPSIS #include <signal.h></p> <p>int sigaction(int <i>signal</i>, const struct sigaction *<i>act</i>, struct sigaction *<i>oldact</i>);</p> <p>DESCRIPTION The sigaction system call is used to change the action taken by a process on receipt of a specific signal. <i>signal</i> specifies the signal and can be any valid signal except SIGKILL and SIGSTOP. If <i>act</i> is non-null, the new action for signal <i>signal</i> is installed from <i>act</i>. If <i>oldact</i> is non-null, the previous action is saved in <i>oldact</i>. The sigaction structure is defined as something like</p> <pre> struct sigaction { void (*<i>sa_handler</i>)(int <i>signal_number</i>); sigset_t <i>sa_mask</i>; int <i>sa_flags</i>; } </pre> <p><i>sa_handler</i> specifies the action to be associated with <i>signal</i> and may be SIG_DFL for the default action, SIG_IGN to ignore this signal, or a pointer to a signal handling function. <i>sa_mask</i> gives a mask of signals which should be blocked during execution of the signal handler. In addition, the signal which triggered the handler will be blocked, unless the SA_NODEFER or SA_NOMASK flags are used. <i>sa_flags</i> specifies a set of flags which modify the behaviour of the signal handling process. It is formed by the bitwise OR of zero or more of the following:</p> <p>SA_NOCLDSTOP If <i>signal</i> is SIGCHLD, do not receive notification when child processes stop (i.e., when child processes receive one of SIGSTOP, SIGTSTP, SIGTTIN or SIGTTOU).</p> <p>SA_RESTART Provide behaviour compatible with BSD signal semantics by making certain system calls restartable across signals. Without SA_RESTART the system calls return an error and set <i>errno</i> to EINTR when interrupted by a signal.</p> <p>RETURN VALUES sigaction(0) returns 0 on success; on error, -1 is returned, and <i>errno</i> is set to indicate the error.</p> <p>ERRORS EINVAL An invalid signal was specified. This will also be generated if an attempt is made to change the action for SIGKILL or SIGSTOP, which cannot be caught.</p> <p>SEE ALSO kill(1), kill(2), killpg(2), pause(2), sigsetops(3).</p>	<p>RETURN VALUE pthread_mutex_init always returns 0. The other mutex functions return 0 on success and a non-zero error code on error.</p> <p>ERRORS The pthread_mutex_lock function returns the following error code on error: EINVAL the mutex has not been properly initialized. EDEADLK the mutex is already locked by the calling thread (“error checking” mutexes only). The pthread_mutex_unlock function returns the following error code on error: EINVAL the mutex has not been properly initialized. EPERM the calling thread does not own the mutex (“error checking” mutexes only). The pthread_mutex_destroy function returns the following error code on error: EBUSY the mutex is currently locked.</p> <p>AUTHOR Xavier Leroy <Xavier.Leroy@inria.fr></p> <p>SEE ALSO pthread_mutexattr_init(3), pthread_mutexattr_setkind_np(3), pthread_cancel(3).</p>	<p>SP-Klausur Manual-Auszug</p> <p>2021-02-24</p> <p>1</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, strncpy, 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.