

**NAME** bind – bind a name to a socket

**SYNOPSIS**

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

int bind(int s, const struct sockaddr *name, int namelen);
```

**DESCRIPTION**

`bind()` assigns a name to an unnamed socket `s`. When a socket is created with `socket(3N)`, it exists in a name space (address family) but has no name assigned. `bind()` requests that the name pointed to by `name` be assigned to the socket.

**RETURN VALUE**

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

**ERRORS**

The `bind()` call will fail if:

- 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** `s` is not a valid descriptor.
- EINVAL** `namelen` 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** `s` is a descriptor for a file, not a socket.

The following errors are specific to binding names in the UNIX domain:

- EACCES** Search permission is denied for a component of the path prefix of the pathname in `name`.
- 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 `name`.
- ENOENT** A component of the path prefix of the pathname in `name` does not exist.
- ENOTDIR** A component of the path prefix of the pathname in `name` is not a directory.
- EROFS** The inode would reside on a read-only file system.

**SEE ALSO** `unlink(2)`, `socket(3N)`, `attributes(5)`, `socket(5)`

**NOTES**

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 `unlink(2)`).

The rules used in name binding vary between communication domains.

**NAME** accept – accept a connection on a socket

**SYNOPSIS**

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

int accept(int s, struct sockaddr *addr, int *addrlen);
```

**DESCRIPTION**

The argument `s` is a socket that has been created with `socket(3N)` and bound to an address with `bind(3N)`, and that is listening for connections after a call to `listen(3N)`. The `accept()` function extracts the first connection on the queue of pending connections, creates a new socket with the properties of `s`, and allocates a new file descriptor, `ns`, for the socket. If no pending connections are present on the queue and the socket is not marked as non-blocking, `accept()` 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, `accept()` returns an error as described below. The `accept()` function uses the `netconfig(4)` file to determine the STREAMS device file name associated with `s`. This is the device on which the connect indication will be accepted. The accepted socket, `ns`, is used to read and write data to and from the socket that connected to `ns`; it is not used to accept more connections. The original socket (`s`) remains open for accepting further connections.

The argument `addr` 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 `addr` parameter is determined by the domain in which the communication occurs.

The argument `addrlen` is a value-result parameter. Initially, it contains the amount of space pointed to by `addr`; on return it contains the length in bytes of the address returned.

The `accept()` function is used with connection-based socket types, currently with `SOCK_STREAM`. It is possible to `select(3C)` or `poll(2)` a socket for the purpose of an `accept()` 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 `accept()`.

**RETURN VALUE**

On success, these system calls return a nonnegative integer that is a file descriptor for the accepted socket. On error, `-1` is returned, and `errno` is set appropriately.

**ERRORS**

`accept()` will fail if:

- 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 `s` could not be found in the `netconfig` 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.
- EWOULDBLOCK** The socket is marked as non-blocking and no connections are present to be accepted.

**SEE ALSO** `poll(2)`, `bind(3N)`, `connect(3N)`, `listen(3N)`, `select(3C)`, `socket(3N)`, `netconfig(4)`, `attributes(5)`, `socket(5)`

fdopen(3) fdopen(3) ipv6/socket(7) ipv6/socket(7)

**NAME**  
fdopen – associate a stream with a file descriptor

**SYNOPSIS**  
#include <stdio.h>  
FILE \*fdopen(int *fdes*, const char \**mode*);

**DESCRIPTION**  
The `fdopen()` function associates a stream with a file descriptor *fdes*, whose value must be less than 255. The *mode* argument is a character string having one of the following values:

```
r or rb      open a file for reading
w or wb      open a file for writing
a or ab      open a file for writing at end of file
r+ or rb+ or r+b  open a file for update (reading and writing)
w+ or wb+ or w+b  open a file for update (reading and writing)
a+ or ab+ or a+b  open a file for update (reading and writing) at end of file
```

The meaning of these flags is exactly as specified in `fopen(3S)`, except that modes beginning with `w` do not cause truncation of the file.

The mode of the stream must be allowed by the file access mode of the open file. The file position indicator associated with the new stream is set to the position indicated by the file offset associated with the file descriptor.

`fdopen()` will preserve the offset maximum previously set for the open file description corresponding to *fdes*.

The error and end-of-file indicators for the stream are cleared. The `fdopen()` function may cause the `st_atime` field of the underlying file to be marked for update.

**RETURN VALUES**  
Upon successful completion, `fdopen()` returns a pointer to a stream. Otherwise, a null pointer is returned and `errno` is set to indicate the error.

`fdopen()` may fail and not set `errno` if there are no free `stdio` streams.

**ERRORS**  
The `fdopen()` function may fail if:

- EBADF** The *fdes* argument is not a valid file descriptor.
- EINVAL** The *mode* argument is not a valid mode.
- EMFILE** `FOPEN_MAX` streams are currently open in the calling process.
- EMFILE** `STREAM_MAX` streams are currently open in the calling process.
- ENOMEM** Insufficient space to allocate a buffer.

**USAGE**  
`STREAM_MAX` is the number of streams that one process can have open at one time. If defined, it has the same value as `FOPEN_MAX`.

File descriptors are obtained from calls like `open(2)`, `dup(2)`, `creat(2)` or `pipe(2)`, which open files but do not return streams. Streams are necessary input for almost all of the Section 3S library routines.

**SEE ALSO**  
`creat(2)`, `dup(2)`, `open(2)`, `pipe(2)`, `fclose(3S)`, `fopen(3S)`, `attributes(5)`

**RETURN VALUES**  
-1 is returned if an error occurs. Otherwise the return value is a descriptor referencing the socket.

**NOTES**  
The `sockaddr_in6` structure is bigger than the generic `sockaddr`. Programs that assume that all address types can be stored safely in a `struct sockaddr` need to be changed to use `struct sockaddr_storage` for that instead.

**SEE ALSO**  
`cmsg(3)`, `ip(7)`

SP-Klausur Manual-Auszug 2022-02-23 1

fdopen(3) fdopen(3) ipv6/socket(7) ipv6/socket(7)

**NAME**  
fdopen – Linux IPv6 protocol implementation

**SYNOPSIS**  
#include <sys/socket.h>  
#include <netinet/in.h>  
  
`tcp6_socket = socket(AF_INET6, SOCK_STREAM, 0);`  
`raw6_socket = socket(AF_INET6, SOCK_RAW, protocol);`  
`udp6_socket = socket(AF_INET6, SOCK_DGRAM, protocol);`

**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 `in6addr_any` variable which has `in6_addr` 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.

**Address Format**

```
struct sockaddr_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 \*/  
};

*sin6\_family* is always set to `AF_INET6`; *sin6\_port* is the protocol port (see *sin\_port* in `ip(7)`); *sin6\_flowinfo* is the IPv6 flow identifier; *sin6\_addr* is the 128-bit IPv6 address; *sin6\_scope\_id* 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 *sin6\_scope\_id* contains the interface index (see `netdevice(7)`)

**RETURN VALUES**  
-1 is returned if an error occurs. Otherwise the return value is a descriptor referencing the socket.

**NOTES**  
The `sockaddr_in6` structure is bigger than the generic `sockaddr`. Programs that assume that all address types can be stored safely in a `struct sockaddr` need to be changed to use `struct sockaddr_storage` for that instead.

**SEE ALSO**  
`cmsg(3)`, `ip(7)`

SP-Klausur Manual-Auszug 2022-02-23 1

```
pthread_cond(3)                                pthread_cond(3)

NAME
pthread_cond_init, pthread_cond_destroy, pthread_cond_signal, pthread_cond_broadcast,
pthread_cond_wait, pthread_cond_timedwait – operations on conditions

SYNOPSIS
#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
*ubstime);
int pthread_cond_destroy(pthread_cond_t *cond);
```

**DESCRIPTION**  
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.

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.

**pthread\_cond\_init** initializes the condition variable *cond*, using the condition attributes specified in *cond\_attr*, or default attributes if *cond\_attr* is **NULL**. The LinuxThreads implementation supports no attributes for conditions, hence the *cond\_attr* parameter is actually ignored.

Variables of type **pthread\_cond\_t** can also be initialized statically, using the constant **PTHREAD\_COND\_INITIALIZER**.

**pthread\_cond\_signal** restarts one or multiple of the threads that are waiting on the condition variable *cond*. If no threads are waiting on *cond*, nothing happens. If several threads are waiting on *cond*, at least one is restarted, but it is not specified which.

**pthread\_cond\_broadcast** restarts all the threads that are waiting on the condition variable *cond*. Nothing happens if no threads are waiting on *cond*.

**pthread\_cond\_wait** atomically unlocks the *mutex* (as per **pthread\_unlock\_mutex**) and waits for the condition variable *cond* to be signaled. The thread execution is suspended and does not consume any CPU time until the condition variable is signaled. The *mutex* must be locked by the calling thread on entrance to **pthread\_cond\_wait**. Before returning to the calling thread, **pthread\_cond\_wait** re-acquires *mutex* (as per **pthread\_lock\_mutex**).

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

```
listen(2)                                     listen(2)

NAME
listen – listen for connections on a socket

SYNOPSIS
#include <sys/types.h> /* See NOTES */
#include <sys/socket.h>

int listen(int sockfd, int backlog);

DESCRIPTION
listen() marks the socket referred to by sockfd as a passive socket, that is, as a socket that will be used to accept incoming connection requests using accept(2).


The sockfd argument is a file descriptor that refers to a socket of type SOCK_STREAM or SOCK_SEQPACKET.



The backlog argument defines the maximum length to which the queue of pending connections for sockfd 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.



RETURN VALUE  

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



ERRORS  

EADDRINUSE  

Another socket is already listening on the same port.



EBADF  

The argument sockfd is not a valid descriptor.



ENOTSOCK  

The argument sockfd is not a socket.



NOTES  

To accept connections, the following steps are performed:



1. A socket is created with socket(2).
2. The socket is bound to a local address using bind(2), so that other sockets may be connect(2)ed to it.
3. A willingness to accept incoming connections and a queue limit for incoming connections are specified with listen().
4. Connections are accepted with accept(2).



If the backlog argument is greater than the value in /proc/sys/net/core/somaxconn, then it is silently truncated to that value; the default value in this file is 128.



EXAMPLE  

See bind(2).



SEE ALSO  

accept(2), bind(2), connect(2), socket(2), socket(7)


```

pthread\_create(pthread\_t \*, const pthread\_attr\_t \*, void \*, void \*)

pthread\_create – create a new thread / pthread\_exit – terminate the calling thread

**NAME**

**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 of control 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\_cond(3)

signaled (and thus ignored) between the time a thread locks the mutex and the time it waits on the condition variable.

**pthread\_cond\_timedwait** atomically unlocks *mutex* and waits on *cond*, as **pthread\_cond\_wait** does, but it also bounds the duration of the wait. If *cond* has not been signaled within the amount of time specified by *abstime*, the mutex *mutex* is re-acquired and **pthread\_cond\_timedwait** returns the error **ETIMEOUT**. The *abstime* parameter specifies an absolute time, with the same origin as **time(2)** and **gettimeofday(2)**: an *abstime* of 0 corresponds to 00:00:00 GMT, January 1, 1970.

**pthread\_cond\_destory** destroys a condition variable, freeing the resources it might hold. No threads must be waiting on the condition variable on entrance to **pthread\_cond\_destory**. In the LinuxThreads implementation, no resources are associated with condition variables, thus **pthread\_cond\_destory** actually does nothing except checking that the condition has no waiting threads.

**CANCELLATION**

**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 *mutex* argument to **pthread\_cond\_wait** and **pthread\_cond\_timedwait**, and finally executes the cancellation. Consequently, cleanup handlers are assured that *mutex* is locked when they are called.

**ASYNC-SIGNAL SAFETY**

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.

**RETURN VALUE**

All condition variable functions return 0 on success and a non-zero error code on error.

**ERRORS**

**pthread\_cond\_init**, **pthread\_cond\_signal**, **pthread\_cond\_broadcast**, and **pthread\_cond\_wait** never return an error code.

The **pthread\_cond\_timedwait** function returns the following error codes on error:

**ETIMEOUT**  
the condition variable was not signaled until the timeout specified by *abstime*

**EINTR**  
**pthread\_cond\_timedwait** was interrupted by a signal

The **pthread\_cond\_destory** function returns the following error code on error:

**EBUSY**  
some threads are currently waiting on *cond*.

**AUTHOR**  
Xavier Leroy <Xavier.Leroy@inria.fr>

**SEE ALSO**  
**pthread\_condattr\_init(3)**, **pthread\_mutex\_lock(3)**, **pthread\_mutex\_unlock(3)**, **gettimeofday(2)**, **nanosleep(2)**.

pthread\_mutex(3) pthread\_mutex(3) pthread\_mutex(3)

performed before the mutex returns to the unlocked state.

**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**.

**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.

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.

**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.

**RETURN VALUE**  
**pthread\_mutex\_init** always returns 0. The other mutex functions return 0 on success and a non-zero error code on error.

**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.

**AUTHOR**  
Xavier Leroy <Xavier.Leroy@inria.fr>

**SEE ALSO**  
**pthread\_mutexattr\_init(3)**, **pthread\_mutexattr\_setkind\_np(3)**, **pthread\_cancel(3)**.

pthread\_mutex(3) pthread\_mutex(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 errorcheckmutex = 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

<p>sigaction(2)</p> <p><b>NAME</b> sigaction – POSIX signal handling functions.</p> <p><b>SYNOPSIS</b> #include &lt;signal.h&gt;</p> <p><b>DESCRIPTION</b> The <b>sigaction</b> 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 <b>SIGKILL</b> and <b>SIGSTOP</b>. 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>.</p> <p>The <b>sigaction</b> structure is defined as something like</p> <pre> struct sigaction {     void (*sa_handler)(int signal_number);     sigset_t sa_mask;     int sa_flags; } </pre> <p><i>sa_handler</i> specifies the action to be associated with <i>signal</i> and may be <b>SIG_DFL</b> for the default action, <b>SIG_IGN</b> to ignore this signal, or a pointer to a signal handling function.</p> <p><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 <b>SA_NODEFER</b> or <b>SA_NOMASK</b> flags are used.</p> <p><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><b>SA_NOCLDSTOP</b> If <i>signal</i> is <b>SIGCHLD</b>, do not receive notification when child processes stop (i.e., when child processes receive one of <b>SIGSTOP</b>, <b>SIGTSTP</b>, <b>SIGTTOU</b>, <b>SIGTTIN</b> or <b>SIGTTOU</b>).</p> <p><b>SA_RESTART</b> Provide behaviour compatible with BSD signal semantics by making certain system calls restartable across signals. Without <b>SA_RESTART</b> the system calls return an error and set <i>errno</i> to <b>EINTR</b> when interrupted by a signal.</p> <p><b>RETURN VALUES</b> <b>sigaction(0)</b> returns 0 on success; on error, <b>-1</b> is returned, and <i>errno</i> is set to indicate the error.</p> <p><b>ERRORS</b> <b>EINVAL</b> An invalid signal was specified. This will also be generated if an attempt is made to change the action for <b>SIGKILL</b> or <b>SIGSTOP</b>, which cannot be caught.</p> <p><b>SEE ALSO</b> <b>kill(1)</b>, <b>kill(2)</b>, <b>killpg(2)</b>, <b>pause(2)</b>, <b>sigsetops(3)</b>.</p>	<p>string(3)</p> <p><b>NAME</b> strcpy, strncpy, strcmp, strcpy, strdup, strlen, strcat, strcmp, strncpy, strstr, strtok – string operations</p> <p><b>SYNOPSIS</b> #include &lt;string.h&gt;</p> <p><b>char *strcat(char *dest, const char *src);</b> Append the string <i>src</i> to the string <i>dest</i>, returning a pointer <i>dest</i>.</p> <p><b>char *strcpy(const char *s, int c);</b> Return a pointer to the first occurrence of the character <i>c</i> in the string <i>s</i>.</p> <p><b>int strcmp(const char *s1, const char *s2);</b> Compare the strings <i>s1</i> with <i>s2</i>. It returns an integer less than, equal to, or greater than zero if <i>s1</i> is found, respectively, to be less than, to match, or be greater than <i>s2</i>.</p> <p><b>char *strncpy(char *dest, const char *src);</b> Copy the string <i>src</i> to <i>dest</i>, returning a pointer to the start of <i>dest</i>.</p> <p><b>char *strdup(const char *s);</b> Return a duplicate of the string <i>s</i> in memory allocated using <b>malloc(3)</b>.</p> <p><b>size_t strlen(const char *s);</b> Return the length of the string <i>s</i>.</p> <p><b>char *strncat(char *dest, const char *src, size_t n);</b> Append at most <i>n</i> characters from the string <i>src</i> to the string <i>dest</i>, returning a pointer to <i>dest</i>.</p> <p><b>int strncmp(const char *s1, const char *s2, size_t n);</b> Compare at most <i>n</i> bytes of the strings <i>s1</i> and <i>s2</i>. It returns an integer less than, equal to, or greater than zero if <i>s1</i> is found, respectively, to be less than, to match, or be greater than <i>s2</i>.</p> <p><b>char *strncpy(char *dest, const char *src, size_t n);</b> Copy at most <i>n</i> bytes from string <i>src</i> to <i>dest</i>, returning a pointer to the start of <i>dest</i>.</p> <p><b>char *strstr(const char *haystack, const char *needle);</b> Find the first occurrence of the substring <i>needle</i> in the string <i>haystack</i>, returning a pointer to the found substring.</p> <p><b>char *strtok(char *s, const char *delim);</b> Extract tokens from the string <i>s</i> that are delimited by one of the bytes in <i>delim</i>.</p> <p><b>DESCRIPTION</b> The string functions perform operations on null-terminated strings.</p>
<p>sigaction(2)</p>	<p>string(3)</p>
<p>SP-Klausur Manual-Auszug</p> <p>2022-02-23</p> <p>1</p>	<p>SP-Klausur Manual-Auszug</p> <p>2022-02-23</p> <p>1</p>