

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	2019-07-30	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	2019-07-30	1

<p>close(2)</p> <p>NAME close – close a file descriptor</p> <p>SYNOPSIS <code>#include <unistd.h></code> <code>int close(int <i>fd</i>);</code></p> <p>DESCRIPTION <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). 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 <code>close()</code> returns zero on success. On error, <code>-1</code> is returned, and <i>errno</i> is set appropriately.</p> <p>ERRORS 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>	<p>dup(2)</p> <p>NAME dup, dup2 – duplicate a file descriptor</p> <p>SYNOPSIS <code>#include <unistd.h></code> <code>int dup(int <i>oldfd</i>);</code> <code>int dup2(int <i>oldfd</i>, int <i>newfd</i>);</code></p> <p>DESCRIPTION <code>dup()</code> and <code>dup2()</code> create a copy of the file descriptor <i>oldfd</i>. <code>dup()</code> uses the lowest-numbered unused descriptor for the new descriptor. <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: * 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>. 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. 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 <code>dup()</code> and <code>dup2()</code> return the new descriptor, or <code>-1</code> if an error occurred (in which case, <i>errno</i> is set appropriately).</p> <p>ERRORS 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>	<p>dup(2)</p>
<p>close(2)</p>	<p>close(2)</p>	<p>dup(2)</p>
<p>SP-Klausur Manual-Auszug</p>	<p>2019-07-30</p>	<p>2019-07-30</p>
<p>1</p>	<p>1</p>	<p>1</p>

feof/ferror/fileno(3)	feof/ferror/fileno(3)	fopen/fdopen/fileno(3)
NAME	clearerr, feof, ferorr, fileno – check and reset stream status	fopen, fdopen, fileno – stream open functions
SYNOPSIS	#include <stdio.h>	#include <stdio.h>
	<pre>void clearerr(FILE *stream); int feof(FILE *stream); int ferorr(FILE *stream); int fileno(FILE *stream);</pre>	<pre>FILE *fopen(const char *path, const char *mode); FILE *fdopen(int fd, const char *mode); int fileno(FILE *stream);</pre>
DESCRIPTION	<p>The function clearerr() clears the end-of-file and error indicators for the stream pointed to by <i>stream</i>. The function feof() tests the end-of-file indicator for the stream pointed to by <i>stream</i>, 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 <i>stream</i>, returning non-zero if it is set. The error indicator can only be reset by the clearerr() function. The function fileno() examines the argument <i>stream</i> and returns its integer descriptor. For non-locking counterparts, see unlocked_stdio(3).</p>	<p>The fopen function opens the file whose name is the string pointed to by <i>path</i> and associates a stream with it. The argument <i>mode</i> 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.</p>
ERRORS	<p>These functions should not fail and do not set the external variable <i>errno</i>. (However, in case fileno() detects that its argument is not a valid stream, it must return -1 and set <i>errno</i> to EBADF.)</p>	<p>The fdopen function associates a stream with the existing file descriptor, <i>fd</i>. The <i>mode</i> 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 <i>fd</i>, 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.</p>
CONFORMING TO	<p>The functions clearerr(), feof(), and ferorr() conform to C89 and C99.</p>	<p>The function fileno() examines the argument <i>stream</i> and returns its integer descriptor.</p> <p>Upon successful completion fopen, fdopen and freopen return a FILE pointer. Otherwise, NULL is returned and the global variable <i>errno</i> is set to indicate the error.</p>
SEE ALSO	<p>open(2), fdopen(3), stdio(3), unlocked_stdio(3)</p>	
ERRORS	EINVAL	EINVAL
RETURN VALUE		<p>The fopen, fdopen and freopen functions may also fail and set <i>errno</i> for any of the errors specified for the routine malloc(3).</p>
SEE ALSO		<p>The fdopen function may also fail and set <i>errno</i> for any of the errors specified for the routine open(2). The fdopen function may also fail and set <i>errno</i> for any of the errors specified for the routine fcntl(2).</p>
SP-Klausur Manual-Auszug	2019-07-30	2019-07-30
SP-Klausur Manual-Auszug	1	1

getc/fgets/putc/fputs(3) getc/fgets/putc/fputs(3)
ip(v6)/socket(7) ip(v6)/socket(7)

NAME
fgetc, fgets, getc, getchar, fputc, fputs, putc, putchar – input and output of characters and strings

SYNOPSIS
#include <stdio.h>

```
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);
```

DESCRIPTION
fgetc() reads the next character from *stream* and returns it as an *unsigned char* cast to an *int*, or **EOF** on end of file or error.

getc() is equivalent to fgetc() except that it may be implemented as a macro which evaluates *stream* more than once.

getchar() is equivalent to getc(stdin).

fgets() reads in at most one less than *size* characters from *stream* and stores them into the buffer pointed to by *s*. 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 *c*, cast to an *unsigned char*, to *stream*.

fputs() writes the string *s* to *stream*, without its terminating null byte ('\0').

putc() is equivalent to fputc() except that it may be implemented as a macro which evaluates *stream* 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 *stdio* library for the same output stream.

RETURN VALUE
fgetc(), getc() and getchar() return the character read as an *unsigned char* cast to an *int* or **EOF** on end of file or error.

fgets() returns *s* 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 *unsigned char* cast to an *int* or **EOF** on error.

fputs() returns a nonnegative number on success, or **EOF** on error.

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), fputc(3), fputwc(3), fputs(3), fseek(3), fwrite(3), gets(3), putchar(3), scanf(3), ungetc(3), unlocked_stdio(3)

NAME
ip(v6), AF_INET6 – 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_** 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 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 */
};
```

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 *sockaddrr_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 sockaddrr_storage* for that instead.

SEE ALSO

[cmsg\(3\)](#), [ip\(7\)](#)

listen(2)	<p>listen – listen for connections on a socket</p> <p>NAME listen – listen for connections on a socket</p> <p>SYNOPSIS</p> <pre>#include <sys/types.h> /* See NOTES */ #include <sys/socket.h></pre> <p>int listen(int sockfd, int backlog);</p> <p>DESCRIPTION</p> <p>listen() marks the socket referred to by <i>sockfd</i> as a passive socket, that is, as a socket that will be used to accept incoming connection requests using accept(2).</p> <p>The <i>sockfd</i> argument is a file descriptor that refers to a socket of type SOCK_STREAM or SOCK_SEQPACKET.</p> <p>The <i>backlog</i> argument defines the maximum length to which the queue of pending connections for <i>sockfd</i> 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> <p>RETURN VALUE</p> <p>On success, zero is returned. On error, -1 is returned, and <i>errno</i> is set appropriately.</p> <p>ERRORS</p> <p>EADDRINUSE Another socket is already listening on the same port.</p> <p>EBADF The argument <i>sockfd</i> is not a valid descriptor.</p> <p>ENOTSOCK The argument <i>sockfd</i> is not a socket.</p> <p>NOTES</p> <p>To accept connections, the following steps are performed:</p> <ol style="list-style-type: none"> 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). <p>If the <i>backlog</i> argument is greater than the value in <i>/proc/sys/net/core/somaxconn</i>, then it is silently truncated to that value; the default value in this file is 128.</p> <p>EXAMPLE</p> <p>See bind(2).</p> <p>SEE ALSO</p> <p>accept(2), bind(2), connect(2), socket(2), socket(7)</p>	listen(2)
printf/sprintf(3)	<p>printf, sprintf, snprintf, vprintf, vsprintf, vsnprintf – formatted output conversion</p> <p>NAME printf, sprintf, snprintf, vprintf, vsprintf, vsnprintf – formatted output conversion</p> <p>SYNOPSIS</p> <pre>#include <stdio.h></pre> <pre>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>DESCRIPTION</p> <p>The functions in the printf() family produce output according to a <i>format</i> as described below. The function printf() writes output to <i>stdout</i>, the standard output stream; fprintf() writes output to the given output stream; sprintf() and snprintf() write to the character string <i>str</i>.</p> <p>The function snprintf() writes at most <i>size</i> bytes (including the trailing null byte $\backslash0$) 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 stdarg(3)) are converted for output.</p> <p>Return value</p> <p>Upon successful return, these functions return the number of characters printed (not including the trailing $\backslash0$ used to end output to strings).</p> <p>The functions sprintf() and vsnprintf() do not write more than <i>size</i> bytes (including the trailing $\backslash0$). If the output was truncated due to this limit then the return value is the number of characters (not including the trailing $\backslash0$) 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>Format of the format string</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>The conversion specifier</p> <p>A character that specifies the type of conversion to be applied. An example for a conversion specifier is:</p> <p>0, u, x, X</p> <p>The <i>unsigned int</i> argument is converted to unsigned octal (o), unsigned decimal (t), or unsigned hexadecimal (x and X) notation.</p> <p>s</p> <p>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 $\backslash0$; 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>SEE ALSO</p> <p>printf(1), asprintf(3), dprintf(3), scanf(3), setlocale(3), wcrtomb(3), wprintf(3), locale(5)</p>	printf/sprintf(3)
SP-Klausur Manual-Auszug	2019-07-30	1

<p>pthread_create(pthread_exit(3))</p> <p>pthread_create – create a new thread / pthread_exit – terminate the calling thread</p> <p>NAME</p> <p>pthread_create / pthread_exit</p> <p>SYNOPSIS</p> <pre>#include <pthread.h> int pthread_create(pthread_t * thread, pthread_attr_t * attr, void * (*start_routine)(void *), void * arg); void pthread_exit(void *retval);</pre> <p>DESCRIPTION</p> <p>pthread_create 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 pthread_exit(3), or implicitly, by returning from the <i>start_routine</i> function. The latter case is equivalent to calling pthread_exit(3) 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 pthread_attr_init(3) for a complete list of thread attributes. The <i>attr</i> 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.</p> <p>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.</p> <p>The <i>retval</i> argument is the return value of the thread. It can be consulted from another thread using pthread_join(3).</p> <p>RETURN VALUE</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 pthread_exit function never returns.</p> <p>ERRORS</p> <p>EAGAIN not enough system resources to create a process for the new thread.</p> <p>EAGAIN more than PTHREAD_THREADS_MAX threads are already active.</p> <p>AUTHOR</p> <p>Xavier Leroy <Xavier.Leroy@inria.fr></p> <p>SEE ALSO</p> <p>pthread_join(3), pthread_detach(3), pthread_attr_init(3),</p>	<p>pthread_join(3)</p> <p>pthread_join – join with a terminated thread</p> <p>NAME</p> <p>pthread_join</p> <p>SYNOPSIS</p> <pre>#include <pthread.h> int pthread_join(pthread_t thread, void **retval);</pre> <p>Compile and link with <i>-pthread</i>.</p> <p>DESCRIPTION</p> <p>The pthread_join(0) function waits for the thread specified by <i>thread</i> to terminate. If that thread has already terminated, then pthread_join(0) returns immediately. The thread specified by <i>thread</i> must be joinable.</p> <p>If <i>retval</i> is not NULL, then pthread_join(0) copies the exit status of the target thread (i.e., the value that the target thread supplied to pthread_exit(3)) into the location pointed to by <i>retval</i>. If the target thread was canceled, then PTHREAD_CANCELED 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 pthread_join(0) is canceled, then the target thread will remain joinable (i.e., it will not be detached).</p> <p>RETURN VALUE</p> <p>On success, pthread_join(0) returns 0; on error, it returns an error number.</p> <p>ERRORS</p> <p>EDEADLK A deadlock was detected (e.g., two threads tried to join with each other); or <i>thread</i> specifies the calling thread.</p> <p>EINVAL <i>thread</i> is not a joinable thread.</p> <p>EINVAL Another thread is already waiting to join with this thread.</p> <p>ESRCH No thread with the ID <i>thread</i> could be found.</p> <p>NOTES</p> <p>After a successful call to pthread_join(0), 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, &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>EXAMPLE</p> <p>See pthread_create(3).</p> <p>SEE ALSO</p> <p>pthread_cancel(3), pthread_create(3), pthread_detach(3), pthread_exit(3), pthread_t(7)</p>	<p>pthread_create(pthread_exit(3))</p> <p>pthread_create – create a new thread / pthread_exit – terminate the calling thread</p> <p>NAME</p> <p>pthread_create / pthread_exit</p> <p>SYNOPSIS</p> <pre>#include <pthread.h> int pthread_create(pthread_t * thread, pthread_attr_t * attr, void * (*start_routine)(void *), void * arg); void pthread_exit(void *retval);</pre> <p>DESCRIPTION</p> <p>pthread_create 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 pthread_exit(3), or implicitly, by returning from the <i>start_routine</i> function. The latter case is equivalent to calling pthread_exit(3) 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 pthread_attr_init(3) for a complete list of thread attributes. The <i>attr</i> 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.</p> <p>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.</p> <p>The <i>retval</i> argument is the return value of the thread. It can be consulted from another thread using pthread_join(3).</p> <p>RETURN VALUE</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 pthread_exit function never returns.</p> <p>ERRORS</p> <p>EAGAIN not enough system resources to create a process for the new thread.</p> <p>EAGAIN more than PTHREAD_THREADS_MAX threads are already active.</p> <p>AUTHOR</p> <p>Xavier Leroy <Xavier.Leroy@inria.fr></p> <p>SEE ALSO</p> <p>pthread_join(3), pthread_detach(3), pthread_attr_init(3),</p>	<p>pthread_join(3)</p> <p>pthread_join – join with a terminated thread</p> <p>NAME</p> <p>pthread_join</p> <p>SYNOPSIS</p> <pre>#include <pthread.h> int pthread_join(pthread_t thread, void **retval);</pre> <p>Compile and link with <i>-pthread</i>.</p> <p>DESCRIPTION</p> <p>The pthread_join(0) function waits for the thread specified by <i>thread</i> to terminate. If that thread has already terminated, then pthread_join(0) returns immediately. The thread specified by <i>thread</i> must be joinable.</p> <p>If <i>retval</i> is not NULL, then pthread_join(0) copies the exit status of the target thread (i.e., the value that the target thread supplied to pthread_exit(3)) into the location pointed to by <i>retval</i>. If the target thread was canceled, then PTHREAD_CANCELED 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 pthread_join(0) is canceled, then the target thread will remain joinable (i.e., it will not be detached).</p> <p>RETURN VALUE</p> <p>On success, pthread_join(0) returns 0; on error, it returns an error number.</p> <p>ERRORS</p> <p>EDEADLK A deadlock was detected (e.g., two threads tried to join with each other); or <i>thread</i> specifies the calling thread.</p> <p>EINVAL <i>thread</i> is not a joinable thread.</p> <p>EINVAL Another thread is already waiting to join with this thread.</p> <p>ESRCH No thread with the ID <i>thread</i> could be found.</p> <p>NOTES</p> <p>After a successful call to pthread_join(0), 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, &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>EXAMPLE</p> <p>See pthread_create(3).</p> <p>SEE ALSO</p> <p>pthread_cancel(3), pthread_create(3), pthread_detach(3), pthread_exit(3), pthread_t(7)</p>	<p>SP-Klausur Manual-Auszug</p> <p>2019-07-30</p> <p>1</p>
---	--	---	--	--

pthread_sigmask(3)	pthread_sigmask(3)	sigaction(2)	sigaction(2)
<p>NAME</p> <p>pthread_sigmask – examine and change mask of blocked signals</p> <p>SYNOPSIS</p> <pre>#include <signal.h> int pthread_sigmask(int how, const sigset_t *set, sigset_t *oset);</pre> <p>DESCRIPTION</p> <p>The <code>pthread_sigmask()</code> function is just like <code>sigprocmask()</code>, with the difference that its use in multi-threaded programs is explicitly specified by POSIX.1.</p> <p>The <code>sigprocmask()</code> function is used to examine and/or change the caller's signal mask. If the value is <code>SIG_BLOCK</code>, the set pointed to by the argument <code>set</code> is added to the current signal mask. If the value is <code>SIG_UNBLOCK</code>, the set pointed by the argument <code>set</code> is removed from the current signal mask. If the value is <code>SIG_SETMASK</code>, the current signal mask is replaced by the set pointed to by the argument <code>set</code>. If the argument <code>oset</code> is not <code>NULL</code>, the previous mask is stored in the space pointed to by <code>oset</code>. If the value of the argument <code>set</code> is <code>NULL</code>, the value <code>how</code> is not significant and the caller's signal mask is unchanged; thus, the call can be used to inquire about currently blocked signals.</p> <p>If there are any pending unblocked signals after the call to <code>sigprocmask()</code>, at least one of those signals will be delivered before the call to <code>sigprocmask()</code> returns.</p> <p>It is not possible to block those signals that cannot be ignored this restriction is silently imposed by the system. See <code>sigaction(2)</code>.</p> <p>If <code>sigprocmask()</code> fails, the caller's signal mask is not changed.</p> <p>RETURN VALUES</p> <p>On success, <code>pthread_sigmask()</code> returns <code>0</code>. On failure, it returns an error number.</p> <p>ERRORS</p> <p><code>pthread_sigmask()</code> fails if any of the following is true:</p> <p>EFAULT <code>set</code> or <code>oset</code> points to an illegal address.</p> <p>EINVAL The value of the <code>how</code> argument is not equal to one of the defined values.</p>	<p>NAME</p> <p>sigaction – POSIX signal handling functions.</p> <p>SYNOPSIS</p> <pre>#include <signal.h> int sigaction(int signal, const struct sigaction *act, struct sigaction *oldact);</pre> <p>DESCRIPTION</p> <p>The <code>sigaction</code> system call is used to change the action taken by a process on receipt of a specific signal. <code>signal</code> specifies the signal and can be any valid signal except <code>SIGKILL</code> and <code>SIGSTOP</code>. If <code>act</code> is non-null, the new action for signal <code>signal</code> is installed from <code>act</code>. If <code>oldact</code> is non-null, the previous action is saved in <code>oldact</code>.</p> <p>The <code>sigaction</code> 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><code>sa_handler</code> specifies the action to be associated with <code>signal</code> and may be <code>SIG_DFL</code> for the default action, <code>SIG_IGN</code> to ignore this signal, or a pointer to a signal handling function.</p> <p><code>sa_mask</code> 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 <code>SA_NODEFER</code> or <code>SA_NOMASK</code> flags are used.</p> <p><code>sa_flags</code> 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 <code>signal</code> is <code>SIGCHLD</code>, do not receive notification when child processes stop (i.e., when child processes receive one of <code>SIGSTOP</code>, <code>SIGTSTP</code>, <code>SIGTTIN</code> or <code>SIGTTOU</code>).</p> <p>SA_RESTART Provide behaviour compatible with BSD signal semantics by making certain system calls restartable across signals. Without <code>SA_RESTART</code> the system calls return an error and set <code>errno</code> to <code>EINTR</code> when interrupted by a signal.</p> <p>RETURN VALUES</p> <p><code>sigaction()</code> returns 0 on success; on error, <code>-1</code> is returned, and <code>errno</code> is set to indicate the error.</p> <p>ERRORS</p> <p>EINVAL An invalid signal was specified. This will also be generated if an attempt is made to change the action for <code>SIGKILL</code> or <code>SIGSTOP</code>, which cannot be caught.</p> <p>SEE ALSO</p> <p><code>kill(1)</code>, <code>kill(2)</code>, <code>killpg(2)</code>, <code>pause(2)</code>, <code>sigsetops(3)</code>,</p>	<p>RETURN VALUES</p> <p><code>sigaction()</code> returns 0 on success; on error, <code>-1</code> is returned, and <code>errno</code> is set to indicate the error.</p> <p>ERRORS</p> <p>EINVAL An invalid signal was specified. This will also be generated if an attempt is made to change the action for <code>SIGKILL</code> or <code>SIGSTOP</code>, which cannot be caught.</p> <p>SEE ALSO</p> <p><code>kill(1)</code>, <code>kill(2)</code>, <code>killpg(2)</code>, <code>pause(2)</code>, <code>sigsetops(3)</code>,</p>	<p>SP-Klausur Manual-Auszug</p> <p>2019-07-30</p> <p>1</p>

sigsetops(3C)	sigsetops(3C)	string(3)	string(3)
NAME	sigsetops, sigemptyset, sigfillset, sigaddset, sigdelset, sigismember – manipulate sets of signals		
SYNOPSIS	<pre>#include <signal.h> int sigemptyset(sigset_t *set); int sigfillset(sigset_t *set); int sigaddset(sigset_t *set, int signo); int sigdelset(sigset_t *set, int signo); int sigismember(sigset_t *set, int signo);</pre>		
DESCRIPTION	<p>These functions manipulate <i>sigset_t</i> data types, representing the set of signals supported by the implementation.</p> <p>sigemptyset() initializes the set pointed to by <i>set</i> to exclude all signals defined by the system.</p> <p>sigfillset() initializes the set pointed to by <i>set</i> to include all signals defined by the system.</p> <p>sigaddset() adds the individual signal specified by the value of <i>signo</i> to the set pointed to by <i>set</i>.</p> <p>sigdelset() deletes the individual signal specified by the value of <i>signo</i> from the set pointed to by <i>set</i>.</p> <p>sigismember() checks whether the signal specified by the value of <i>signo</i> is a member of the set pointed to by <i>set</i>.</p> <p>Any object of type <i>sigset_t</i> must be initialized by applying either sigemptyset() or sigfillset() before applying any other operation.</p>		
RETURN VALUES	Upon successful completion, the sigismember() function returns a value of one if the specified signal is a member of the specified set, or a value of 0 if it is not. Upon successful completion, the other functions return a value of 0. Otherwise a value of -1 is returned and errno is set to indicate the error.		
ERRORS	<p>sigaddset(), sigdelset(), and sigismember() will fail if the following is true:</p> <p>EINVAL The value of the <i>signo</i> argument is not a valid signal number.</p> <p>sigfillset() will fail if the following is true:</p> <p>EFAULT The <i>set</i> argument specifies an invalid address.</p>		
SEE ALSO	sigaction(2) , sigpending(2) , sigprocmask(2) , sigsuspend(2) , attributes(5) , signal(5)		
SP-Klausur Manual-Auszug	2019-07-30	SP-Klausur Manual-Auszug	2019-07-30
1	1	1	1