

accept(2)

accept(2)

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 *s*; 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 VALUES

The `accept()` function returns `-1` on error. If it succeeds, it returns a non-negative integer that is a descriptor for the accepted socket.

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 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.

SEE ALSO

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

bind(2)

bind(2)

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. 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 VALUES

If the bind is successful, `0` is returned. A return value of `-1` indicates an error, which is further specified in the global `errno`.

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.

ENOSOCK *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.

dup(2)

dup(2)

NAME

dup, dup2 – duplicate a file descriptor

SYNOPSIS

```
#include <unistd.h>
```

```
int dup(int oldfd);  
int dup2(int oldfd, int newfd);
```

DESCRIPTION

dup() and dup2() create a copy of the file descriptor *oldfd*.

dup() uses the lowest-numbered unused descriptor for the new descriptor.

dup2() makes *newfd* be the copy of *oldfd*, closing *newfd* first if necessary, but note the following:

- * If *oldfd* is not a valid file descriptor, then the call fails, and *newfd* is not closed.
- * If *oldfd* is a valid file descriptor, and *newfd* has the same value as *oldfd*, then dup2() does nothing, and returns *newfd*.

After a successful return from dup() or dup2(), the old and new file descriptors may be used interchangeably. They refer to the same open file description (see open(2)) and thus share file offset and file status flags; for example, if the file offset is modified by using lseek(2) 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 (FD_CLOEXEC; seefcntl(2)) for the duplicate descriptor is off.

RETURN VALUE

dup() and dup2() return the new descriptor, or -1 if an error occurred (in which case, *errno* is set appropriately).

ERRORS

EBADF

oldfd isn't an open file descriptor, or *newfd* is out of the allowed range for file descriptors.

EBUSY

(Linux only) This may be returned by dup2() during a race condition with open(2) and dup().

EINTR

The dup2() call was interrupted by a signal; see signal(7).

EMFILE

The process already has the maximum number of file descriptors open and tried to open a new one.

NOTES

The error returned by dup2() is different from that returned by fcntl(..., F_DUPFD, ...) when *newfd* is out of range. On some systems dup2() also sometimes returns EINVAL like F_DUPFD.

If *newfd* was open, any errors that would have been reported at close(2) time are lost. A careful programmer will not use dup2() without closing *newfd* first.

SEE ALSO

close(2), fcntl(2), open(2)

feof/ferro/fileno(3)

feof/ferro/fileno(3)

NAME

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

SYNOPSIS

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

```
void clearerr(FILE *stream);  
int feof(FILE *stream);  
int ferro(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 ferro() 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 ferro() conform to C89 and C99.

SEE ALSO

open(2), filopen(3), stdio(3), unlocked_stdio(3)

fdopen/fdopen/filen0(3)

open/fdopen/filen0(3)

NAME

fdopen, fdopen, fileno – stream open functions

SYNOPSIS

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

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

DESCRIPTION

The **fdopen** 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'd, 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.

RETURN VALUE

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

ERRORS

EINVAL

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

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

The **fdopen** 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)**.

SEE ALSO

open(2), **fclose(3)**, **fileno(3)**

socket(2) / ipv6(7)

socket(2) / ipv6(7)

NAME

ipv6, PF_INET6 – Linux IPv6 protocol implementation

SYNOPSIS

```
#include <sys/socket.h>
#include <netinet/in.h>
```

```
tcp6_socket = socket(PF_INET6, SOCK_STREAM, 0);
raw6_socket = socket(PF_INET6, SOCK_RAW, IPPROTO);
udp6_socket = socket(PF_INET6, SOCK_DGRAM, IPPROTO);
```

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 on 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 **netdevic(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)**

listen(2)

listen(2)

NAME

listen – listen for connections on a socket

SYNOPSIS

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

```
int listenfd, 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_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 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_detach(3)`

`pthread_detach(3)`

`sigaction(2)`

`sigaction(2)`

NAME

`pthread_detach` – put a running thread in the detached state

SYNOPSIS

```
#include <pthread.h>

int pthread_detach(pthread_t th);
```

DESCRIPTION

pthread_detach put the thread *th* in the detached state. This guarantees that the memory resources consumed by *th* will be freed immediately when *th* terminates. However, this prevents other threads from synchronizing on the termination of *th* using **pthread_join**.

A thread can be created initially in the detached state, using the **detachstate** attribute to **pthread_create(3)**. In contrast, **pthread_detach** applies to threads created in the joinable state, and which need to be put in the detached state later.

After **pthread_detach** completes, subsequent attempts to perform **pthread_join** on *th* will fail. If another thread is already joining the thread *th* at the time **pthread_detach** is called, **pthread_detach** does nothing and leaves *th* in the joinable state.

RETURN VALUE

On success, 0 is returned. On error, a non-zero error code is returned.

ERRORS

ESRCH

No thread could be found corresponding to that specified by *th*

EINVAL

the thread *th* is already in the detached state

AUTHOR

Xavier Leroy <XavierLeroy@inria.fr>

SEE ALSO

pthread_create(3), **pthread_join(3)**, **pthread_attr_setdetachstate(3)**.

NAME

`sigaction` – POSIX signal handling functions.

SYNOPSIS

```
#include <signal.h>

int sigaction(int signum, const struct sigaction *act, struct sigaction *oldact);
```

DESCRIPTION

The **sigaction** system call is used to change the action taken by a process on receipt of a specific signal. *signum* specifies the signal and can be any valid signal except **SIGKILL** and **SIGSTOP**. If *act* is non-null, the new action for signal *signum* is installed from *act*. If *oldact* is non-null, the previous action is saved in *oldact*. The **sigaction** structure is defined as something like

```
struct sigaction {
    void (*sa_handler)(int);
    void (*sa_sigaction)(int, siginfo_t *, void *);
    sigset_t sa_mask;
    int sa_flags;
    void (*sa_restorer)(void);
}
```

On some architectures a union is involved – do not assign to both *sa_handler* and *sa_sigaction*. The *sa_restorer* element is obsolete and should not be used. POSIX does not specify a *sa_restorer* element.

sa_handler specifies the action to be associated with *signum* and may be **SIG_DFL** for the default action, **SIG_IGN** to ignore this signal, or a pointer to a signal handling function.

sa_mask 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.

sa_flags 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:

SA_NOCLDSTOP

If *signum* is **SIGCHLD**, do not receive notification when child processes stop (i.e., when child processes receive one of **SIGSTOP**, **SIGTSTP**, **SIGTTIN** or **SIGTTOU**).

SA_RESTART

Provide behaviour compatible with BSD signal semantics by making certain system calls restartable across signals.

RETURN VALUES

sigaction returns 0 on success and -1 on error.

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.

SEE ALSO

kill(1), **kill(2)**, **killpg(2)**, **pause(2)**, **sigsetops(3)**.

sigsetops(3C)

sigsetops(3C)

printf(3)

printf(3)

NAME

sigsetops, sigemptyset, sigfillset, sigaddset, sigdelset, sigismember – manipulate sets of signals

SYNOPSIS

```
#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);
```

DESCRIPTION

These functions manipulate *sigset_t* data types, representing the set of signals supported by the implementation.

sigemptyset() initializes the set pointed to by *set* to exclude all signals defined by the system.

sigfillset() initializes the set pointed to by *set* to include all signals defined by the system.

sigaddset() adds the individual signal specified by the value of *signo* to the set pointed to by *set*.

sigdelset() deletes the individual signal specified by the value of *signo* from the set pointed to by *set*.

sigismember() checks whether the signal specified by the value of *signo* is a member of the set pointed to by *set*.

Any object of type *sigset_t* must be initialized by applying either **sigemptyset()** or **sigfillset()** before applying any other operation.

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

sigaddset(), **sigdelset()**, and **sigismember()** will fail if the following is true:

EINVAL The value of the *signo* argument is not a valid signal number.

sigfillset() will fail if the following is true:

EFAULT The *set* argument specifies an invalid address.

SEE ALSO

sigaction(2), **sigpending(2)**, **sigprocmask(2)**, **sigsuspend(2)**, **attributes(5)**, **signal(5)**

NAME

printf, fprintf, sprintf, snprintf, vprintf, vfprintf, vsprintf, vsnprintf – formatted output conversion

SYNOPSIS

```
#include <stdio.h>

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, ...);
...
```

DESCRIPTION

The functions in the **printf()** family produce output according to a *format* as described below. The functions **printf()** and **vprintf()** write output to *stdout*, the standard output stream; **fprintf()** and **vfprintf()** write output to the given output stream; **sprintf()**, **snprintf()**, **vsprintf()** and **vsnprintf()** write to the character string *str*.

The functions **sprintf()** and **vsnprintf()** write at most *size* bytes (including the trailing null byte '\0') to *str*.

These eight functions write the output under the control of a *format* string that specifies how subsequent arguments (or arguments accessed via the variable-length argument facilities of **stdarg(3)**) are converted for output.

Return value

Upon successful return, these functions return the number of characters printed (not including the trailing '\0' used to end output to strings).

The functions **sprintf()** and **vsnprintf()** do not write more than *size* bytes (including the trailing '\0'). If the output was truncated due to this limit then the return value is the number of characters (not including the trailing '\0') which would have been written to the final string if enough space had been available. Thus, a return value of *size* or more means that the output was truncated. (See also below under **NOTES**.)

If an output error is encountered, a negative value is returned.

Format of the format string

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 *conversion specifier*. In between there may be (in this order) *zero* or more *flags*, an optional minimum *field width*, an optional *precision* and an optional *length modifier*.

The conversion specifier

A character that specifies the type of conversion to be applied. An example for a conversion specifier is:

s The *conv char ** argument is expected to be a pointer to an array of character type (pointer to a string). Characters from the array are written up to (but not including) a terminating null byte ('\0'); if a precision is specified, no more than the number specified are written. If a precision is given, no null byte need be present; if the precision is not specified, or is greater than the size of the array, the array must contain a terminating null byte.

SEE ALSO

printf(1), **asprintf(3)**, **dprintf(3)**, **scanf(3)**, **setlocale(3)**, **wctomb(3)**, **wprintf(3)**, **locale(5)**

strcmp(3)

strcmp(3)

NAME

strcmp, strcmp – compare two strings

SYNOPSIS

#include <string.h>

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

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

DESCRIPTION

The **strcmp()** function compares the two 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*.

The **strcmp()** function is similar, except it only compares the first (at most) *n* characters of *s1* and *s2*.

RETURN VALUE

The **strcmp()** and **strcmp()** functions return an integer less than, equal to, or greater than zero if *s1* (or the first *n* bytes thereof) is found, respectively, to be less than, to match, or be greater than *s2*.

CONFORMING TO

SVr4, 4.3BSD, C89, C99.

SEE ALSO

bcmp(3), **memcmp(3)**, **strncasecmp(3)**, **strcoll(3)**, **strncasecmp(3)**, **wscmp(3)**, **wscnmp(3)**

strtok(3)

strtok(3)

NAME

strtok, strtok_r – extract tokens from strings

SYNOPSIS

#include <string.h>

char *strtok(char *str, const char *delim);

char *strtok_r(char *str, const char *delim, char **saveptr);

DESCRIPTION

The **strtok()** function breaks a string into a sequence of zero or more nonempty tokens. On the first call to **strtok()** the string to be parsed should be specified in *str*. In each subsequent call that should parse the same string, *str* must be NULL.

The *delim* argument specifies a set of bytes that delimit the tokens in the parsed string. The caller may specify different strings in *delim* in successive calls that parse the same string.

Each call to **strtok()** returns a pointer to a null-terminated string containing the next token. This string does not include the delimiting byte. If no more tokens are found, **strtok()** returns NULL.

A sequence of calls to **strtok()** that operate on the same string maintains a pointer that determines the point from which to start searching for the next token. The first call to **strtok()** sets this pointer to point to the first byte of the string. The start of the next token is determined by scanning forward for the next nondelimiter byte in *str*. If such a byte is found, it is taken as the start of the next token. If no such byte is found, then there are no more tokens, and **strtok()** returns NULL. (A string that is empty or that contains only delimiters will thus cause **strtok()** to return NULL on the first call.)

The end of each token is found by scanning forward until either the next delimiter byte is found or until the terminating null byte ('\0') is encountered. If a delimiter byte is found, it is overwritten with a null byte to terminate the current token, and **strtok()** saves a pointer to the following byte; that pointer will be used as the starting point when searching for the next token. In this case, **strtok()** returns a pointer to the start of the found token.

From the above description, it follows that a sequence of two or more contiguous delimiter bytes in the parsed string is considered to be a single delimiter, and that delimiter bytes at the start or end of the string are ignored. Put another way: the tokens returned by **strtok()** are always nonempty strings. Thus, for example, given the string "aaa;bbb", successive calls to **strtok()** that specify the delimiter string ";" would return the strings "aaa" and "bbb", and then a null pointer.

The **strtok_r()** function is a reentrant version **strtok()**. The *saveptr* argument is a pointer to a *char** variable that is used internally by **strtok_r()** in order to maintain context between successive calls that parse the same string. On the first call to **strtok_r()**, *str* should point to the string to be parsed, and the value of *saveptr* is ignored. In subsequent calls, *str* should be NULL, and *saveptr* should be unchanged since the previous call.

Different strings may be parsed concurrently using sequences of calls to **strtok_r()** that specify different *saveptr* arguments.

RETURN VALUE

strtok() and **strtok_r()** return a pointer to the next token, or NULL, if there are no more tokens.

ATTRIBUTES

Multithreading (see pthreads(7))

The **strtok()** function is not thread-safe, the **strtok_r()** function is thread-safe.

qsort(3)

qsort(3)

qsort(3)

qsort(3)

NAME

qsort, qsort_r – sort an array

SYNOPSIS

```
#include <stdlib.h>
```

```
void qsort(void *base, size_t membh, size_t size,
           int (*compar)(const void *, const void *));
```

```
void qsort_r(void *base, size_t membh, size_t size,
             int (*compar)(const void *, const void *, void *),
             void *arg);
```

Feature Test Macro Requirements for glibc (see [feature_test_macros\(7\)](#)):

```
qsort_r(): _GNU_SOURCE
```

DESCRIPTION

The **qsort()** function sorts an array with *membh* elements of size *size*. The *base* argument points to the start of the array.

The contents of the array are sorted in ascending order according to a comparison function pointed to by *compar*, which is called with two arguments that point to the objects being compared.

The comparison function must return an integer less than, equal to, or greater than zero if the first argument is considered to be respectively less than, equal to, or greater than the second. If two members compare as equal, their order in the sorted array is undefined.

The **qsort_r()** function is identical to **qsort()** except that the comparison function *compar* takes a third argument. A pointer is passed to the comparison function via *arg*. In this way, the comparison function does not need to use global variables to pass through arbitrary arguments, and is therefore reentrant and safe to use in threads.

If no global variables are needed in the comparison function *compar*, **qsort()** is also safe to use in threads.

RETURN VALUE

The **qsort()** and **qsort_r()** functions return no value.

VERSIONS

qsort_r() was added to glibc in version 2.8.

CONFORMING TO

The **qsort()** function conforms to SVr4, 4.3BSD, C89, C99.

NOTES

Library routines suitable for use as the *compar* argument to **qsort()** include [alphasort\(3\)](#) and [version-sort\(3\)](#). To compare C strings, the comparison function can call [strcmp\(3\)](#), as shown in the example below.

EXAMPLE

For one example of use, see the example under [bsearch\(3\)](#).

Another example is the following program, which sorts the strings given in its command-line arguments:

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

static int
cmpstring(const void *p1, const void *p2)
{
```

```
/* The actual arguments to this function are "pointers to
pointers to char", but strcmp(3) arguments are "pointers
to char", hence the following cast plus dereference */
return strcmp((char * const *) p1, *(char * const *) p2);
}
int
main(int argc, char *argv[])
{
    int j;
    if (argc < 2) {
        fprintf(stderr, "Usage: %s <string>... \n", argv[0]);
        exit(EXIT_FAILURE);
    }
    qsort(&argv[1], argc - 1, sizeof(char *), cmpstring);
    for (j = 1; j < argc; j++)
        puts(argv[j]);
    exit(EXIT_SUCCESS);
}

```

SEE ALSO

[sort\(1\)](#), [alphasort\(3\)](#), [strcmp\(3\)](#), [versionsort\(3\)](#)

COLOPHON

This page is part of release 3.74 of the Linux *man-pages* project. A description of the project, information about reporting bugs, and the latest version of this page, can be found at <http://www.kernel.org/doc/man-pages/>.