

**NAME**

intro – introduction to system services and error numbers

**SYNOPSIS**

```
#include <errno.h>
```

**DESCRIPTION**

This section describes all of the system calls.

A 2V section number means one or more of the following:

- The man page documents System V behavior only.
- The man page documents default SunOS behavior and System V behavior as it differs from the default behavior. These System V differences are presented under SYSTEM V section headers.
- The man page documents behavior compliant with *IEEE Std 1003.1-1988* (POSIX.1).

Compile programs for the System V environment using `/usr/5bin/cc`. Compile programs for the default SunOS environment using `/usr/bin/cc`. The following man pages describe the various environments provided by Sun: `lint(1V)`, `ansic(7V)`, `bsd(7)`, `posix(7V)`, `sunos(7V)`, `svidii(7V)`, `svidiii(7V)`, `xopen(7V)`.

Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible return value. This is almost always '-1'; the individual descriptions specify the details. An error code is also made available in the external variable `errno`. `errno` is not cleared on successful calls, so it should be tested only after an error has been indicated. Note: several system calls overload the meanings of these error numbers, and the meanings must be interpreted according to the type and circumstances of the call. See **ERROR CODES** below for a list of system error codes.

As with normal arguments, all return codes and values from functions are of type integer unless otherwise noted.

The rest of this man page is organized as follows:

<b>SYSTEM PARAMETERS</b>	System limits, values and options.
<b>DEFINITIONS</b>	System abstractions and services.
<b>STREAMS</b>	Modular communication between software layers (tty system, networking).
<b>SYSTEM V IPC</b>	System V shared memory, semaphores, and messages.
<b>ERROR CODES</b>	A list of system error codes with descriptions.
<b>LIST OF SYSTEM CALLS</b>	A list of all system calls with brief descriptions.

**SYSTEM PARAMETERS**

Sections 2 and 3 support a naming convention for those system parameters that may change from one object to another (for example, path name length may be 255 on a UFS file system but may be 14 on an NFS file system exported by a System V based server). Typically, the system has to be queried (using `pathconf(2V)`, `fpathconf()`, or `sysconf(2V)`) to retrieve the parameter of interest. The parameters have conceptual names such as `PATH_MAX`. These names are defined in header files if and only if they are invariant across all file systems and releases of the operating system, that is, very rarely. Because they *may* be defined and/or available from the system calls, there have to be separate names for the parameters and their values. The notation `{PATH_MAX}` denotes the value of the parameter `PATH_MAX`. Do not confuse this with `_PC_PATH_MAX`, the name that is passed to the system call to retrieve the value:

```
maxpathlen = pathconf(".", _PC_PATH_MAX);
```

See `pathconf(2V)`, and `sysconf(2V)` for further information about these parameters.

## DEFINITIONS

**Controlling Terminal**

A terminal that is associated with a session. Each session may have at most one controlling terminal; a terminal may be the controlling terminal of at most one session. The controlling terminal is used to direct signals (such as interrupts and job control signals) to the appropriate processes by way of the tty's process group. Controlling terminals are assigned when a session leader opens a terminal file that is not currently a controlling terminal.

**Descriptor**

An integer assigned by the system when a file is referenced by `open(2V)`, `dup(2V)`, or `pipe(2V)` or a socket is referenced by `socket(2)` or `socketpair(2)` that uniquely identifies an access path to that file or socket from a given process or any of its children.

**Directory**

A directory is a special type of file that contains entries that are references to other files. Directory entries are called links. By convention, a directory contains at least two links, `'.'` and `'..'`, referred to as *dot* and *dot-dot* respectively. Dot refers to the directory itself and dot-dot refers to its parent directory.

**Effective User ID, Effective Group ID, and Access Groups**

Access to system resources is governed by three values: the effective user ID, the effective group ID, and the supplementary group ID.

The effective user ID and effective group ID are initially the process's real user ID and real group ID respectively. Either may be modified through execution of a set-user-ID or set-group-ID file (possibly by one of its ancestors) (see `execve(2V)`).

The supplementary group ID are an additional set of group ID's used only in determining resource accessibility. Access checks are performed as described below in **File Access Permissions**.

**File Access Permissions**

Every file in the file system has a set of access permissions. These permissions are used in determining whether a process may perform a requested operation on the file (such as opening a file for writing). Access permissions are established at the time a file is created. They may be changed at some later time through the `chmod(2V)` call.

File access is broken down according to whether a file may be: read, written, or executed. Directory files use the execute permission to control if the directory may be searched.

File access permissions are interpreted by the system as they apply to three different classes of users: the owner of the file, those users in the file's group, anyone else. Every file has an independent set of access permissions for each of these classes. When an access check is made, the system decides if permission should be granted by checking the access information applicable to the caller.

Read, write, and execute/search permissions on a file are granted to a process if:

The process's effective user ID is that of the super-user.

The process's effective user ID matches the user ID of the owner of the file and the owner permissions allow the access.

The process's effective user ID does not match the user ID of the owner of the file, and either the process's effective group ID matches the group ID of the file, or the group ID of the file is in the process's supplementary group IDs, and the group permissions allow the access.

Neither the effective user ID nor effective group ID and supplementary group IDs of the process match the corresponding user ID and group ID of the file, but the permissions for "other users" allow access.

Otherwise, permission is denied.

**File Name**

Names consisting of up to {NAME\_MAX} characters may be used to name an ordinary file, special file, or directory.

These characters may be selected from the set of all ASCII character excluding \0 (null) and the ASCII code for / (slash). (The parity bit, bit 8, must be 0.)

Note: it is generally unwise to use \*, ?, [, or ] as part of file names because of the special meaning attached to these characters by the shell. See `sh(1)`. Although permitted, it is advisable to avoid the use of unprintable characters in file names.

**Parent Process ID**

A new process is created by a currently active process `fork(2V)`. The parent process ID of a process is the process ID of its creator.

**Path Name and Path Prefix**

A path name is a null-terminated character string starting with an optional slash (/), followed by zero or more directory names separated by slashes, optionally followed by a file name. The total length of a path name must be less than {PATH\_MAX} characters.

More precisely, a path name is a null-terminated character string constructed as follows:

```
<path-name>::=<file-name>| <path-prefix><file-name>| /
<path-prefix>::=<rtprefix>| /<rtprefix>
<rtprefix>::=<dirname>| /<rtprefix><dirname>|
```

where `<file-name>` is a string of 1 to {NAME\_MAX} characters other than the ASCII slash and null, and `<dirname>` is a string of 1 to {NAME\_MAX} characters (other than the ASCII slash and null) that names a directory.

If a path name begins with a slash, the search begins at the *root* directory. Otherwise, the search begins at the current working directory.

A slash, by itself, names the root directory. A dot (.) names the current working directory.

A null path name also refers to the current directory. However, this is not true of all UNIX systems. (On such systems, accidental use of a null path name in routines that do not check for it may corrupt the current working directory.) For portable code, specify the current directory explicitly using `"."`, rather than `""`.

**Process Group ID**

Each active process is a member of a process group that is identified by a positive integer called the process group ID. This ID is the process ID of the group leader. This grouping permits the signaling of related processes (see the description of `killpg()` on `kill(2V)`) and the job control mechanisms of `csh(1)`. Process groups exist from their creation until the last member is reaped (that is, a parent issued a call to `wait(2V)`).

**Process ID**

Each active process in the system is uniquely identified by a positive integer called a process ID. The range of this ID is from 0 to MAXPID (see `<sys/param.h>`).

**Real User ID and Real Group ID**

Each user on the system is identified by a positive integer termed the real user ID.

Each user is also a member of one or more groups. One of these groups is distinguished from others and used in implementing accounting facilities. The positive integer corresponding to this distinguished group is termed the real group ID.

All processes have a real user ID and real group ID. These are initialized from the equivalent attributes of the process that created it.

**Root Directory and Current Working Directory**

Each process has associated with it a concept of a root directory and a current working directory for the purpose of resolving path name searches. The root directory is used as the starting point for absolute path name resolution. The current working directory is used as the starting point for relative path name resolution. A process's root directory need not be (but typically is) the root directory of the root file system.

**Session**

Each process is a member of a session. A session is associated with each controlling terminal in the system, such as login shells and windows. Each process is created in the session of its parent. A process may alter its session using `setsid(2V)` if it is not already a session leader. The system supports session IDs. A session leader is a process having process ID equal to process group ID equal to session ID. Only a session leader may acquire a controlling terminal. In SunOS Release 4.1, processes are created in sessions by `init(8)` and `inetd(8C)`. Sessions are also created for processes that disassociate themselves from a controlling terminal using

```
ioctl(fd, TIOCNOTTY, 0)
```

or

```
setpgrp(mygid, 0) For more information about sessions, see setsid(2V).
```

**Signal**

Signals are used for notification of asynchronous events. Signals may be directed to processes, process groups, and other combinations of processes. Signals may be sent by a process or by the operating system. Some signals may be caught. There is typically a default behavior on receipt if they are not caught. For more information about signals, see `signal(3V)`, `kill(2V)`, `sigvec(2)`, `termio(4)`.

**Sockets and Address Families**

A socket is an endpoint for communication between processes, similar to the way a telephone is the endpoint of communication between humans. Each socket has queues for sending and receiving data.

Sockets are typed according to their communications properties. These properties include whether messages sent and received at a socket require the name of the partner, whether communication is reliable, the format used in naming message recipients, etc.

Each instance of the system supports some collection of socket types; consult `socket(2)` for more information about the types available and their properties.

Each instance of the system supports some number of sets of communications protocols. Each protocol set supports addresses of a certain format. An Address Family is the set of addresses for a specific group of protocols. Each socket has an address chosen from the address family in which the socket was created.

**Special Processes**

The processes with a process ID's of 0, 1, and 2 are special. Process 0 is the scheduler. Process 1 is the initialization process `init`, and is the ancestor of every other process in the system. It is used to control the process structure. Process 2 is the paging daemon.

**Super-user**

A process is recognized as a *super-user* process and is granted special privileges if its effective user ID is 0.

**Tty Process Group**

Each active process can be a member of a terminal group that is identified by a positive integer called the tty process group ID. This grouping is used to arbitrate between multiple jobs contending for the same terminal (see `csh(1)`, and `termio(4)`), to direct signals (tty and job control) to the appropriate process group, and to terminate a group of related processes upon termination of one of the processes in the group (see `exit(2V)` and `sigvec(2)`).

**STREAMS**

A set of kernel mechanisms that support the development of network services and data communication *drivers*. It defines interface standards for character input/output within the kernel and between the kernel and user level processes. The STREAMS mechanism is composed of utility routines, kernel facilities and a set of data structures.

**Stream**

A stream is a full-duplex data path within the kernel between a user process and driver routines. The primary components are a stream head, a *driver* and zero or more *modules* between the stream head and *driver*. A stream is analogous to a Shell pipeline except that data flow and processing are bidirectional.

**Stream Head**

In a stream, the stream head is the end of the stream that provides the interface between the stream and a user process. The principle functions of the stream head are processing STREAMS-related system calls, and passing data and information between a user process and the stream.

**Driver**

In a stream, the *driver* provides the interface between peripheral hardware and the stream. A *driver* can also be a pseudo-*driver*, such as a *multiplexor* or *emulator*, and need not be associated with a hardware device.

**Module**

A module is an entity containing processing routines for input and output data. It always exists in the middle of a stream, between the stream's head and a *driver*. A *module* is the STREAMS counterpart to the commands in a Shell pipeline except that a module contains a pair of functions which allow independent bidirectional (*downstream* and *upstream*) data flow and processing.

**Downstream**

In a stream, the direction from stream head to *driver*.

**Upstream**

In a stream, the direction from *driver* to stream head.

**Message**

In a stream, one or more blocks of data or information, with associated STREAMS control structures. Messages can be of several defined types, which identify the message contents. Messages are the only means of transferring data and communicating within a stream.

**Message Queue**

In a stream, a linked list of *messages* awaiting processing by a *module* or *driver*.

**Read Queue**

In a stream, the *message queue* in a *module* or *driver* containing *messages* moving *upstream*.

**Write Queue**

In a stream, the *message queue* in a *module* or *driver* containing *messages* moving *downstream*.

**Multiplexor**

A multiplexor is a driver that allows STREAMS associated with several user processes to be connected to a single *driver*, or several *drivers* to be connected to a single user process. STREAMS does not provide a general multiplexing *driver*, but does provide the facilities for constructing them, and for connecting multiplexed configurations of STREAMS.

**SYSTEM V IPC**

The SunOS system supports the System V IPC namespace. For information about shared memory, semaphores and messages see `msgctl(2)`, `msgget(2)`, `msgop(2)`, `semctl(2)`, `semget(2)`, `semop(2)`, `shmctl(2)`, `shmget(2)` and `shmop(2)`.

**ERROR CODES**

Each system call description attempts to list all possible error numbers. The following is a complete list of the error numbers and their names as given in `<errno.h>`.

**E2BIG 7 Arg list too long**

An argument list longer than 1,048,576 bytes is presented to `execve(2V)` or a routine that called `execve()`.

**EACCES 13 Permission denied**

An attempt was made to access a file in a way forbidden by the protection system.

**EADDRINUSE 48 Address already in use**

Only one usage of each address is normally permitted.

**EADDRNOTAVAIL 49 Can't assign requested address**

Normally results from an attempt to create a socket with an address not on this machine.

**EADV 83 Advertise error**

An attempt was made to advertise a resource which has been advertised already, or to stop the RFS while there are resources still advertised, or to force unmount a resource when it is still advertised. This error is RFS specific.

**EAFNOSUPPORT 47 Address family not supported by protocol family**

An address incompatible with the requested protocol was used. For example, you should not necessarily expect to be able to use PUP Internet addresses with ARPA Internet protocols.

**EAGAIN 11 No more processes**

A `fork(2V)` failed because the system's process table is full or the user is not allowed to create any more processes, or a system call failed because of insufficient resources.

**EALREADY 37 Operation already in progress**

An operation was attempted on a non-blocking object that already had an operation in progress.

**EBADF 9 Bad file number**

Either a file descriptor refers to no open file, or a read (respectively, write) request is made to a file that is open only for writing (respectively, reading).

**EBADMSG 76 Not a data message**

During a `read(2V)`, `getmsg(2)`, or `ioctl(2) I_RECVFD` system call to a STREAMS device, something has come to the head of the queue that cannot be processed. That something depends on the system call

`read(2V)` control information or a passed file descriptor.

`getmsg(2)` passed file descriptor.

`ioctl(2)` control or data information.

**EBUSY 16 Device busy**

An attempt was made to mount a file system that was already mounted or an attempt was made to dismount a file system on which there is an active file (open file, mapped file, current directory, or mounted-on directory).

**ECHILD 10 No children**

A `wait(2V)` was executed by a process that had no existing or unwaited-for child processes.

**ECOMM 85 Communication error on send**

An attempt was made to send messages to a remote machine when no virtual circuit could be found. This error is RFS specific.

**ECONNABORTED 53 Software caused connection abort**

A connection abort was caused internal to your host machine.

**ECONNREFUSED 61 Connection refused**

No connection could be made because the target machine actively refused it. This usually results from trying to connect to a service that is inactive on the foreign host.

- ECONNRESET 54** Connection reset by peer  
A connection was forcibly closed by a peer. This normally results from the peer executing a **shutdown(2)** call.
- EDEADLK 78** Deadlock situation detected/avoided  
An attempt was made to lock a system resource that would have resulted in a deadlock situation.
- EDESTADDRREQ 39** Destination address required  
A required address was omitted from an operation on a socket.
- EDOM 33** Math argument  
The argument of a function in the math library (as described in section 3M) is out of the domain of the function.
- EDQUOT 69** Disc quota exceeded  
A **write()** to an ordinary file, the creation of a directory or symbolic link, or the creation of a directory entry failed because the user's quota of disk blocks was exhausted, or the allocation of an inode for a newly created file failed because the user's quota of inodes was exhausted.
- EEXIST 17** File exists  
An existing file was mentioned in an inappropriate context, for example, **link(2V)**.
- EFAULT 14** Bad address  
The system encountered a hardware fault in attempting to access the arguments of a system call.
- EFBIG 27** File too large  
The size of a file exceeded the maximum file size (1,082,201,088 bytes).
- EHOSTDOWN 64** Host is down  
A socket operation failed because the destination host was down.
- EHOSTUNREACH 65** Host is unreachable  
A socket operation was attempted to an unreachable host.
- EIDRM 77** Identifier removed  
This error is returned to processes that resume execution due to the removal of an identifier.
- EINPROGRESS 36** Operation now in progress  
An operation that takes a long time to complete (such as a **connect(2)**) was attempted on a non-blocking object (see **ioctl(2)**).
- EINTR 4** Interrupted system call  
An asynchronous signal (such as **interrupt** or **quit**) that the process has elected to catch occurred during a system call. If execution is resumed after processing the signal, and the system call is not restarted, it will appear as if the interrupted system call returned this error condition.
- EINVAL 22** Invalid argument  
A system call was made with an invalid argument; for example, dismounting a non-mounted file system, mentioning an unknown signal in **sigvec()** or **kill()**, reading or writing a file for which **lseek()** has generated a negative pointer, or some other argument inappropriate for the call. Also set by math functions, see **intro(3)**.
- EIO 5** I/O error  
Some physical I/O error occurred. This error may in some cases occur on a call following the one to which it actually applies.
- EISCONN 56** Socket is already connected  
A **connect()** request was made on an already connected socket; or, a **sendto()** or **sendmsg()** request on a connected socket specified a destination other than the connected party.
- EISDIR 21** Is a directory  
An attempt was made to write on a directory.

- EISDIR 21** Is a directory  
An attempt was made to write on a directory.
- ELOOP 62** Too many levels of symbolic links  
A path name lookup involved more than 20 symbolic links.
- EMFILE 24** Too many open files  
A process tried to have more open files than the system allows a process to have. The customary configuration limit is 64 per process.
- EMLINK 31** Too many links  
An attempt was made to make more than 32767 hard links to a file.
- EMSGSIZE 40** Message too long  
A message sent on a socket was larger than the internal message buffer.
- EMULTIHOP 87** Multihop attempted  
An attempt was made to access remote resources which are not directly accessible. This error is RFS specific.
- ENAMETOOLONG 63** File name too long  
A component of a path name exceeded 255 characters, or an entire path name exceeded 1024 characters.
- ENETDOWN 50** Network is down  
A socket operation encountered a dead network.
- ENETRESET 52** Network dropped connection on reset  
The host you were connected to crashed and rebooted.
- ENETUNREACH 51** Network is unreachable  
A socket operation was attempted to an unreachable network.
- ENFILE 23** File table overflow  
The system's table of open files is full, and temporarily no more `open()` calls can be accepted.
- ENOBUFS 55** No buffer space available  
An operation on a socket or pipe was not performed because the system lacked sufficient buffer space.
- ENODEV 19** No such device  
An attempt was made to apply an inappropriate system call to a device (for example, an attempt to read a write-only device) or an attempt was made to use a device not configured by the system.
- ENOENT 2** No such file or directory  
This error occurs when a file name is specified and the file should exist but does not, or when one of the directories in a path name does not exist.
- ENOEXEC 8** Exec format error  
A request is made to execute a file which, although it has the appropriate permissions, does not start with a valid magic number (see `a.out(5)`).
- ENOLCK 79** No locks available  
A system-imposed limit on the number of simultaneous file and record locks was reached and no more were available at that time.
- ENOLINK 82** Link has be severed  
The link (virtual circuit) connecting to a remote machine is gone. This error is RFS specific.



- ENOMEM 12** Not enough memory  
During an `execve(2V)`, `sbrk()`, or `brk(2)`, a program asks for more address space or swap space than the system is able to supply, or a process size limit would be exceeded. A lack of swap space is normally a temporary condition; however, a lack of address space is not a temporary condition. The maximum size of the text, data, and stack segments is a system parameter. Soft limits may be increased to their corresponding hard limits.
- ENOMSG 75** No message of desired type  
An attempt was made to receive a message of a type that does not exist on the specified message queue; see `msgop(2)`.
- ENONET 80** Machine is not on the network  
A attempt was made to advertise, unadvertise, mount, or unmount remote resources while the machine has not done the proper startup to connect to the network. This error is Remote File Sharing (RFS) specific.
- ENOPROTOPT 42** Option not supported by protocol  
A bad option was specified in a `setsockopt()` or `getsockopt(2)` call.
- ENOSPC 28** No space left on device  
A `write()` to an ordinary file, the creation of a directory or symbolic link, or the creation of a directory entry failed because no more disk blocks are available on the file system, or the allocation of an inode for a newly created file failed because no more inodes are available on the file system.
- ENOSR 74** Out of stream resources  
During a STREAMS `open(2V)`, either no STREAMS queues or no STREAMS head data structures were available.
- ENOSTR 72** Not a stream device  
A `putmsg(2)` or `getmsg(2)` system call was attempted on a file descriptor that is not a STREAMS device.
- ENOSYS 90** Function not implemented  
An attempt was made to use a function that is not available in this implementation.
- ENOTBLK 15** Block device required  
A file that is not a block device was mentioned where a block device was required, for example, in `mount(2V)`.
- ENOTCONN 57** Socket is not connected  
An request to send or receive data was disallowed because the socket is not connected.
- ENOTDIR 20** Not a directory  
A non-directory was specified where a directory is required, for example, in a path prefix or as an argument to `chdir(2V)`.
- ENOTEMPTY 66** Directory not empty  
An attempt was made to remove a directory with entries other than `'&.'` and `'&|.'` by performing a `rmdir()` system call or a `rename()` system call with that directory specified as the target directory.
- ENOTSOCK 38** Socket operation on non-socket  
Self-explanatory.
- ENOTTY 25** Inappropriate ioctl for device  
The code used in an `ioctl()` call is not supported by the object that the file descriptor in the call refers to.
- ENXIO 6** No such device or address  
I/O on a special file refers to a subdevice that does not exist, or beyond the limits of the device. It may also occur when, for example, a tape drive is not on-line or no disk pack is loaded on a drive.

- EOPNOTSUPP 45** Operation not supported on socket  
For example, trying to *accept* a connection on a datagram socket.
- EPERM 1** Not owner  
Typically this error indicates an attempt to modify a file in some way forbidden except to its owner or super-user. It is also returned for attempts by ordinary users to do things allowed only to the super-user.
- EPFNOSUPPORT 46** Protocol family not supported  
The protocol family has not been configured into the system or no implementation for it exists.
- EPIPE 32** Broken pipe  
An attempt was made to write on a pipe or socket for which there is no process to read the data. This condition normally generates a signal; the error is returned if the signal is caught or ignored.
- EPROTO 86** Protocol error  
Some protocol error occurred. This error is device specific, but is generally not related to a hardware failure.
- EPROTONOSUPPORT 43** Protocol not supported  
The protocol has not been configured into the system or no implementation for it exists.
- EPROTOTYPE 41** Protocol wrong type for socket  
A protocol was specified that does not support the semantics of the socket type requested. For example, you cannot use the ARPA Internet UDP protocol with type **SOCK\_STREAM**.
- ERANGE 34** Result too large  
The value of a function in the math library (as described in section 3M) is unrepresentable within machine precision.
- EREMOTE 71** Too many levels of remote in path  
An attempt was made to remotely mount a file system into a path that already has a remotely mounted component.
- EROFS 30** Read-only file system  
An attempt to modify a file or directory was made on a file system mounted read-only.
- ERREMOTE 81** Object is remote  
An attempt was made to advertise a resource which is not on the local machine, or to mount/unmount a device (or pathname) that is on a remote machine. This error is RFS specific.
- ESHUTDOWN 58** Can't send after socket shutdown  
A request to send data was disallowed because the socket had already been shut down with a previous **shutdown(2)** call.
- ESOCKTNOSUPPORT 44** Socket type not supported  
The support for the socket type has not been configured into the system or no implementation for it exists.
- ESPIPE 29** Illegal seek  
An **lseek()** was issued to a socket or pipe. This error may also be issued for other non-seekable devices.
- ESRCH 3** No such process  
The process or process group whose number was given does not exist, or any such process is already dead.
- ESRMNT 84** Srmount error  
An attempt was made to stop RFS while there are resources still mounted by remote machines. This error is RFS specific.

**ESTALE 70** Stale NFS file handle

An NFS client referenced a file that it had opened but that had since been deleted.

**ETIME 73** Timer expired

The timer set for a STREAMS `ioctl(2)` call has expired. The cause of this error is device specific and could indicate either a hardware or software failure, or perhaps a timeout value that is too short for the specific operation. The status of the `ioctl(2)` operation is indeterminate.

**ETIMEDOUT 60** Connection timed out

A `connect` request or an NFS request failed because the party to which the request was made did not properly respond after a period of time. (The timeout period is dependent on the communication protocol.)

**ETXTBSY 26** Text file busy

An attempt was made to execute a pure-procedure program that is currently open for writing, or an attempt was made to open for writing a pure-procedure program that is being executed.

**EUSERS 68** Too many users

An operation to read disk quota information for the user failed because the system quota table was full.

**EWOULDBLOCK 35** Operation would block

An operation that would cause a process to block was attempted on an object in non-blocking mode (see `ioctl(2)`).

**EXDEV 18** Cross-device link

A hard link to a file on another file system was attempted.

unused 0

**SEE ALSO**

`brk(2)`, `chdir(2V)`, `chmod(2V)`, `connect(2)`, `dup(2V)`, `execve(2V)`, `exit(2V)`, `fork(2V)`, `getmsg(2)`, `getsockopt(2)`, `ioctl(2)`, `killpg(2)`, `link(2V)`, `mount(2V)`, `msgctl(2)`, `msgget(2)`, `msgop(2)`, `open(2V)`, `pipe(2V)`, `putmsg(2)`, `read(2V)`, `semctl(2)`, `semget(2)`, `semop(2)`, `getsockopt(2)`, `shmctl(2)`, `shmget(2)`, `shmop(2)`, `shutdown(2)`, `sigvec(2)`, `socket(2)`, `socketpair(2)`, `wait(2V)`, `cs(1)`, `sh(1)`, `intro(3)`, `perror(3)`, `termio(4)`, `a.out(5)`

**LIST OF SYSTEM CALLS****Name Appears on Page Description**

<b>accept</b>	<b>accept(2)</b>	accept a connection on a socket
<b>access</b>	<b>access(2V)</b>	determine accessibility of file
<b>acct</b>	<b>acct(2V)</b>	turn accounting on or off
<b>adjtime</b>	<b>adjtime(2)</b>	correct the time to allow synchronization of the system clock
<b>async_daemon</b>	<b>nfssvc(2)</b>	NFS daemons
<b>audit</b>	<b>audit(2)</b>	write a record to the audit log
<b>auditon</b>	<b>auditon(2)</b>	manipulate auditing
<b>auditsvc</b>	<b>auditsvc(2)</b>	write audit records to specified file descriptor
<b>bind</b>	<b>bind(2)</b>	bind a name to a socket
<b>brk</b>	<b>brk(2)</b>	change data segment size
<b>chdir</b>	<b>chdir(2V)</b>	change current working directory
<b>chmod</b>	<b>chmod(2V)</b>	change mode of file
<b>chown</b>	<b>chown(2V)</b>	change owner and group of a file
<b>chroot</b>	<b>chroot(2)</b>	change root directory
<b>close</b>	<b>close(2V)</b>	delete a descriptor
<b>connect</b>	<b>connect(2)</b>	initiate a connection on a socket
<b>creat</b>	<b>creat(2V)</b>	create a new file
<b>dup</b>	<b>dup(2V)</b>	duplicate a descriptor
<b>dup2</b>	<b>dup(2V)</b>	duplicate a descriptor

<b>execve</b>	<b>execve(2V)</b>	execute a file
<b>_exit</b>	<b>exit(2V)</b>	terminate a process
<b>fchmod</b>	<b>chmod(2V)</b>	change mode of file
<b>fchown</b>	<b>chown(2V)</b>	change owner and group of a file
<b>fcntl</b>	<b>fcntl(2V)</b>	file control
<b>flock</b>	<b>flock(2)</b>	apply or remove an advisory lock on an open file
<b>fork</b>	<b>fork(2V)</b>	create a new process
<b>fpathconf</b>	<b>pathconf(2V)</b>	query file system related limits and options
<b>fstat</b>	<b>stat(2V)</b>	get file status
<b>fstatfs</b>	<b>statfs(2)</b>	get file system statistics
<b>fsync</b>	<b>fsync(2)</b>	synchronize a file's in-core state with that on disk
<b>ftruncate</b>	<b>truncate(2)</b>	set a file to a specified length
<b>getauid</b>	<b>getauid(2)</b>	get and set user audit identity
<b>getdents</b>	<b>getdents(2)</b>	gets directory entries in a filesystem independent format
<b>getdirenties</b>	<b>getdirenties(2)</b>	gets directory entries in a filesystem independent format
<b>getdomainname</b>	<b>getdomainname(2)</b>	get/set name of current domain
<b>getdtablesize</b>	<b>getdtablesize(2)</b>	get descriptor table size
<b>getegid</b>	<b>getgid(2V)</b>	get group identity
<b>geteuid</b>	<b>getuid(2V)</b>	get user identity
<b>getgid</b>	<b>getgid(2V)</b>	get group identity
<b>getgroups</b>	<b>getgroups(2V)</b>	get or set supplementary group IDs
<b>gethostid</b>	<b>gethostid(2)</b>	get unique identifier of current host
<b>gethostname</b>	<b>gethostname(2)</b>	get/set name of current host
<b>getitimer</b>	<b>getitimer(2)</b>	get/set value of interval timer
<b>getmsg</b>	<b>getmsg(2)</b>	get next message from a stream
<b>getpagesize</b>	<b>getpagesize(2)</b>	get system page size
<b>getpeername</b>	<b>getpeername(2)</b>	get name of connected peer
<b>getpgrp</b>	<b>getpgrp(2V)</b>	return or set the process group of a process
<b>getpid</b>	<b>getpid(2V)</b>	get process identification
<b>getppid</b>	<b>getpid(2V)</b>	get process identification
<b>getpriority</b>	<b>getpriority(2)</b>	get/set process nice value
<b>getrlimit</b>	<b>getrlimit(2)</b>	control maximum system resource consumption
<b>getrusage</b>	<b>getrusage(2)</b>	get information about resource utilization
<b>getsockname</b>	<b>getsockname(2)</b>	get socket name
<b>getsockopt</b>	<b>getsockopt(2)</b>	get and set options on sockets
<b>gettimeofday</b>	<b>gettimeofday(2)</b>	get or set the date and time
<b>getuid</b>	<b>getuid(2V)</b>	get user identity
<b>ioctl</b>	<b>ioctl(2)</b>	control device
<b>kill</b>	<b>kill(2V)</b>	send a signal to a process or a group of processes
<b>killpg</b>	<b>killpg(2)</b>	send signal to a process group
<b>link</b>	<b>link(2V)</b>	make a hard link to a file
<b>listen</b>	<b>listen(2)</b>	listen for connections on a socket
<b>lseek</b>	<b>lseek(2V)</b>	move read/write pointer
<b>lstat</b>	<b>stat(2V)</b>	get file status
<b>mctl</b>	<b>mctl(2)</b>	memory management control
<b>mincore</b>	<b>mincore(2)</b>	determine residency of memory pages
<b>mkdir</b>	<b>mkdir(2V)</b>	make a directory file
<b>mkfifo</b>	<b>mknod(2V)</b>	make a special file
<b>mknod</b>	<b>mknod(2V)</b>	make a special file
<b>mmap</b>	<b>mmap(2)</b>	map pages of memory
<b>mount</b>	<b>mount(2V)</b>	mount file system
<b>mprotect</b>	<b>mprotect(2)</b>	set protection of memory mapping
<b>msgctl</b>	<b>msgctl(2)</b>	message control operations

<b>msgget</b>	<b>msgget(2)</b>	get message queue
<b>msgop</b>	<b>msgop(2)</b>	message operations
<b>msgrcv</b>	<b>msgop(2)</b>	message operations
<b>msgsnd</b>	<b>msgop(2)</b>	message operations
<b>msync</b>	<b>msync(2)</b>	synchronize memory with physical storage
<b>munmap</b>	<b>munmap(2)</b>	unmap pages of memory.
<b>nfssvc</b>	<b>nfssvc(2)</b>	NFS daemons
<b>open</b>	<b>open(2V)</b>	open or create a file for reading or writing
<b>pathconf</b>	<b>pathconf(2V)</b>	query file system related limits and options
<b>pipe</b>	<b>pipe(2V)</b>	create an interprocess communication channel
<b>poll</b>	<b>poll(2)</b>	I/O multiplexing
<b>profil</b>	<b>profil(2)</b>	execution time profile
<b>ptrace</b>	<b>ptrace(2)</b>	process trace
<b>putmsg</b>	<b>putmsg(2)</b>	send a message on a stream
<b>quotactl</b>	<b>quotactl(2)</b>	manipulate disk quotas
<b>read</b>	<b>read(2V)</b>	read input
<b>readlink</b>	<b>readlink(2)</b>	read value of a symbolic link
<b>readv</b>	<b>read(2V)</b>	read input
<b>reboot</b>	<b>reboot(2)</b>	reboot system or halt processor
<b>recv</b>	<b>recv(2)</b>	receive a message from a socket
<b>recvfrom</b>	<b>recv(2)</b>	receive a message from a socket
<b>recvmsg</b>	<b>recv(2)</b>	receive a message from a socket
<b>rename</b>	<b>rename(2V)</b>	change the name of a file
<b>rmdir</b>	<b>rmdir(2V)</b>	remove a directory file
<b>sbrk</b>	<b>brk(2)</b>	change data segment size
<b>select</b>	<b>select(2)</b>	synchronous I/O multiplexing
<b>semctl</b>	<b>semctl(2)</b>	semaphore control operations
<b>semget</b>	<b>semget(2)</b>	get set of semaphores
<b>semop</b>	<b>semop(2)</b>	semaphore operations
<b>send</b>	<b>send(2)</b>	send a message from a socket
<b>sendmsg</b>	<b>send(2)</b>	send a message from a socket
<b>sendto</b>	<b>send(2)</b>	send a message from a socket
<b>setaudit</b>	<b>setuseraudit(2)</b>	set the audit classes for a specified user ID
<b>setaudit</b>	<b>getaudit(2)</b>	get and set user audit identity
<b>setdomainname</b>	<b>getdomainname(2)</b>	get/set name of current domain
<b>setgroups</b>	<b>getgroups(2V)</b>	get or set supplementary group IDs
<b>sethostname</b>	<b>gethostname(2)</b>	get/set name of current host
<b>setitimer</b>	<b>getitimer(2)</b>	get/set value of interval timer
<b>setpgid</b>	<b>setpgid(2V)</b>	set process group ID for job control
<b>setpgrp</b>	<b>getpgrp(2V)</b>	return or set the process group of a process
<b>setpriority</b>	<b>getpriority(2)</b>	get/set process nice value
<b>setregid</b>	<b>setregid(2)</b>	set real and effective group IDs
<b>setreuid</b>	<b>setreuid(2)</b>	set real and effective user IDs
<b>setrlimit</b>	<b>getrlimit(2)</b>	control maximum system resource consumption
<b>setsid</b>	<b>setsid(2V)</b>	create session and set process group ID
<b>setsockopt</b>	<b>getsockopt(2)</b>	get and set options on sockets
<b>settimeofday</b>	<b>gettimeofday(2)</b>	get or set the date and time
<b>setuseraudit</b>	<b>setuseraudit(2)</b>	set the audit classes for a specified user ID
<b>sgetl</b>	<b>sputl(2)</b>	access long integer data in a machine-independent fashion
<b>shmat</b>	<b>shmop(2)</b>	shared memory operations
<b>shmctl</b>	<b>shmctl(2)</b>	shared memory control operations
<b>shmdt</b>	<b>shmop(2)</b>	shared memory operations
<b>shmget</b>	<b>shmget(2)</b>	get shared memory segment identifier

<b>shmop</b>	<b>shmop(2)</b>	shared memory operations
<b>shutdown</b>	<b>shutdown(2)</b>	shut down part of a full-duplex connection
<b>sigblock</b>	<b>sigblock(2)</b>	block signals
<b>sigmask</b>	<b>sigblock(2)</b>	block signals
<b>sigpause</b>	<b>sigpause(2V)</b>	automatically release blocked signals and wait for interrupt
<b>sigpending</b>	<b>sigpending(2V)</b>	examine pending signals
<b>sigprocmask</b>	<b>sigprocmask(2V)</b>	examine and change blocked signals
<b>sigsetmask</b>	<b>sigsetmask(2)</b>	set current signal mask
<b>sigstack</b>	<b>sigstack(2)</b>	set and/or get signal stack context
<b>sigsuspend</b>	<b>sigpause(2V)</b>	automatically release blocked signals and wait for interrupt
<b>sigvec</b>	<b>sigvec(2)</b>	software signal facilities
<b>socket</b>	<b>socket(2)</b>	create an endpoint for communication
<b>socketpair</b>	<b>socketpair(2)</b>	create a pair of connected sockets
<b>sputl</b>	<b>sputl(2)</b>	access long integer data in a machine-independent fashion
<b>stat</b>	<b>stat(2V)</b>	get file status
<b>statfs</b>	<b>statfs(2)</b>	get file system statistics
<b>swapon</b>	<b>swapon(2)</b>	add a swap device for interleaved paging/swapping
<b>symlink</b>	<b>symlink(2)</b>	make symbolic link to a file
<b>sync</b>	<b>sync(2)</b>	update super-block
<b>syscall</b>	<b>syscall(2)</b>	indirect system call
<b>sysconf</b>	<b>sysconf(2V)</b>	query system related limits, values, options
<b>tell</b>	<b>lseek(2V)</b>	move read/write pointer
<b>truncate</b>	<b>truncate(2)</b>	set a file to a specified length
<b>umask</b>	<b>umask(2V)</b>	set file creation mode mask
<b>umount</b>	<b>umount(2V)</b>	remove a file system
<b>uname</b>	<b>uname(2V)</b>	get information about current system
<b>unlink</b>	<b>unlink(2V)</b>	remove directory entry
<b>unmount</b>	<b>umount(2V)</b>	remove a file system
<b>ustat</b>	<b>ustat(2)</b>	get file system statistics
<b>utimes</b>	<b>utimes(2)</b>	set file times
<b>vadvise</b>	<b>vadvise(2)</b>	give advice to paging system
<b>vfork</b>	<b>vfork(2)</b>	spawn new process in a virtual memory efficient way
<b>vhangup</b>	<b>vhangup(2)</b>	virtually "hangup" the current control terminal
<b>wait</b>	<b>wait(2V)</b>	wait for process to terminate or stop, examine returned status
<b>wait3</b>	<b>wait(2V)</b>	wait for process to terminate or stop, examine returned status
<b>wait4</b>	<b>wait(2V)</b>	wait for process to terminate or stop, examine returned status
<b>waitpid</b>	<b>wait(2V)</b>	wait for process to terminate or stop, examine returned status
<b>WEXITSTATUS</b>	<b>wait(2V)</b>	wait for process to terminate or stop, examine returned status
<b>WIFEXITED</b>	<b>wait(2V)</b>	wait for process to terminate or stop, examine returned status
<b>WIFSIGNALED</b>	<b>wait(2V)</b>	wait for process to terminate or stop, examine returned status
<b>WIFSTOPPED</b>	<b>wait(2V)</b>	wait for process to terminate or stop, examine returned status
<b>write</b>	<b>write(2V)</b>	write output
<b>writev</b>	<b>write(2V)</b>	write output
<b>WSTOPSIG</b>	<b>wait(2V)</b>	wait for process to terminate or stop, examine returned status
<b>WTERMSIG</b>	<b>wait(2V)</b>	wait for process to terminate or stop, examine returned status

**NAME**

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

**SYNOPSIS**

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

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

**DESCRIPTION**

The argument *s* is a socket that has been created with **socket(2)**, bound to an address with **bind(2)**, and is listening for connections after a **listen(2)**. **accept()** extracts the first connection on the queue of pending connections, creates a new socket with the same properties of *s* and allocates a new file descriptor 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 non-blocking and no pending connections are present on the queue, **accept()** returns an error as described below. The accepted socket is used to read and write data to and from the socket which connected to this one; 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 known to the communications layer. The exact format of the *addr* parameter is determined by the domain in which the communication is occurring. The *addrlen* is a value-result parameter; it should initially contain the amount of space pointed to by *addr*; on return it will contain the actual length (in bytes) of the address returned. This call is used with connection-based socket types, currently with **SOCK\_STREAM**.

It is possible to **select(2)** a socket for the purposes of doing an **accept()** by selecting it for read.

**RETURN VALUES**

**accept()** returns a non-negative descriptor for the accepted socket on success. On failure, it returns **-1** and sets **errno** to indicate the error.

**ERRORS**

<b>EBADF</b>	The descriptor is invalid.
<b>EFAULT</b>	The <i>addr</i> parameter is not in a writable part of the user address space.
<b>ENOTSOCK</b>	The descriptor references a file, not a socket.
<b>EOPNOTSUPP</b>	The referenced socket is not of type <b>SOCK_STREAM</b> .
<b>EWOULDBLOCK</b>	The socket is marked non-blocking and no connections are present to be accepted.

**SEE ALSO**

**bind(2)**, **connect(2)**, **listen(2)**, **select(2)**, **socket(2)**

**NAME**

access – determine accessibility of file

**SYNOPSIS**

```
#include <unistd.h>

int access(path, mode)
char *path;
int mode;
```

**DESCRIPTION**

*path* points to a path name naming a file. `access()` checks the named file for accessibility according to *mode*, which is an inclusive or of the following bits:

<b>R_OK</b>	test for read permission
<b>W_OK</b>	test for write permission
<b>X_OK</b>	test for execute or search permission

The following value may also be supplied for *mode*:

<b>F_OK</b>	test whether the directories leading to the file can be searched and the file exists.
-------------	---

The real user ID and the supplementary group IDs (including the real group ID) are used in verifying permission, so this call is useful to set-UID programs.

Notice that only access bits are checked. A directory may be indicated as writable by `access()`, but an attempt to open it for writing will fail (although files may be created there); a file may look executable, but `execve()` will fail unless it is in proper format.

**RETURN VALUES**

`access()` returns:

0	on success.
-1	on failure and sets <code>errno</code> to indicate the error.

**ERRORS**

<b>EACCES</b>	Search permission is denied for a component of the path prefix of <i>path</i> . The file access permissions do not permit the requested access to the file named by <i>path</i> .
<b>EFAULT</b>	<i>path</i> points outside the process's allocated address space.
<b>EINVAL</b>	An invalid value was specified for <i>mode</i> .
<b>EIO</b>	An I/O error occurred while reading from or writing to the file system.
<b>ELOOP</b>	Too many symbolic links were encountered in translating <i>path</i> .
<b>ENAMETOOLONG</b>	The length of the path argument exceeds <code>{PATH_MAX}</code> . A pathname component is longer than <code>{NAME_MAX}</code> while <code>{_POSIX_NO_TRUNC}</code> is in effect (see <code>pathconf(2V)</code> ).
<b>ENOENT</b>	The file named by <i>path</i> does not exist.
<b>ENOTDIR</b>	A component of the path prefix of <i>path</i> is not a directory.
<b>EROFS</b>	The file named by <i>path</i> is on a read-only file system and write access was requested.

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

<b>ENOENT</b>	<i>path</i> points to an empty string.
---------------	--



**SEE ALSO**

**chmod(2V), stat(2V)**

**NAME**

`acct` – turn accounting on or off

**SYNOPSIS**

```
int acct (path)
char *path;
```

**DESCRIPTION**

`acct()` is used to enable or disable the process accounting. If process accounting is enabled, an accounting record will be written on an accounting file for each process that terminates. Termination can be caused by one of two things: an `exit()` call or a signal; see `exit(2V)` and `sigvec(2)`. The effective user ID of the calling process must be super-user to use this call.

*path* points to a path name naming the accounting file. The accounting file format is given in `acct(5)`.

The accounting routine is enabled if *path* is not a NULL pointer and no errors occur during the system call. It is disabled if *path* is a NULL pointer and no errors occur during the system call.

If accounting is already turned on, and a successful `acct()` call is made with a non-NULL *path*, all subsequent accounting records will be written to the new accounting file.

**SYSTEM V DESCRIPTION**

If accounting is already turned on, it is an error to call `acct()` with a non-NULL *path*.

**RETURN VALUES**

`acct()` returns:

- 0        on success.
- 1       on failure and sets `errno` to indicate the error.

**ERRORS**

EACCES	Search permission is denied for a component of the path prefix of <i>path</i> . The file referred to by <i>path</i> is not a regular file.
EFAULT	<i>path</i> points outside the process's allocated address space.
EINVAL	Support for accounting was not configured into the system.
EIO	An I/O error occurred while reading from or writing to the file system.
ELOOP	Too many symbolic links were encountered in translating the path name.
ENAMETOOLONG	The length of the path argument exceeds <code>{PATH_MAX}</code> . A pathname component is longer than <code>{NAME_MAX}</code> (see <code>sysconf(2V)</code> ) while <code>{_POSIX_NO_TRUNC}</code> is in effect (see <code>pathconf(2V)</code> ).
ENOENT	The named file does not exist.
ENOTDIR	A component of the path prefix of <i>path</i> is not a directory.
EPERM	The caller is not the super-user.
EROFS	The named file resides on a read-only file system.

**SYSTEM V ERRORS**

EBUSY	<i>path</i> is non-NULL, and accounting is already turned on.
ENOENT	<i>path</i> points to an empty string.

**SEE ALSO**

`exit(2V)`, `sigvec(2)`, `acct(5)`, `sa(8)`

**BUGS**

No accounting records are produced for programs running when a crash occurs. In particular non-terminating programs are never accounted for.

**NOTES**

Accounting is automatically disabled when free space on the file system the accounting file resides on drops below 2 percent; it is enabled when free space rises above 4 percent.

**NAME**

**adjtime** – correct the time to allow synchronization of the system clock

**SYNOPSIS**

```
#include <sys/time.h>

int adjtime(delta, olddelta)
struct timeval *delta;
struct timeval *olddelta;
```

**DESCRIPTION**

**adjtime()** adjusts the system's notion of the current time, as returned by **gettimeofday(2)**, advancing or retarding it by the amount of time specified in the **struct timeval** (defined in **<sys/time.h>**) pointed to by *delta*.

The adjustment is effected by speeding up (if that amount of time is positive) or slowing down (if that amount of time is negative) the system's clock by some small percentage, generally a fraction of one percent. Thus, the time is always a monotonically increasing function. A time correction from an earlier call to **adjtime()** may not be finished when **adjtime()** is called again. If *olddelta* is not a NULL pointer, then the structure it points to will contain, upon return, the number of microseconds still to be corrected from the earlier call. If *olddelta* is a NULL pointer, the corresponding information will not be returned.

This call may be used in time servers that synchronize the clocks of computers in a local area network. Such time servers would slow down the clocks of some machines and speed up the clocks of others to bring them to the average network time.

Only the super-user may adjust the time of day.

The adjustment value will be silently rounded to the resolution of the system clock.

**RETURN**

A 0 return value indicates that the call succeeded. A -1 return value indicates an error occurred, and in this case an error code is stored into the global variable **errno**.

**ERRORS**

<b>EFAULT</b>	<i>delta</i> or <i>olddelta</i> points outside the process's allocated address space.
	<i>olddelta</i> points to a region of the process' allocated address space that is not writable.
<b>EPERM</b>	The process's effective user ID is not that of the super-user.

**SEE ALSO**

**date(1V)**, **gettimeofday(2)**

**NAME**

audit – write a record to the audit log

**SYNOPSIS**

```
#include <sys/label.h>
#include <sys/audit.h>

int audit (record)
audit_record_t *record;
```

**DESCRIPTION**

The `audit()` system call is used to write a record to the system audit log file. The data pointed to by *record* is written to the audit log file. The data should be a well-formed audit record as described by `audit.log(5)`. The kernel sets the time stamp value in the record and performs a minimal check on the data before writing it to the audit log file.

Only the super-user may successfully execute this call.

**RETURN VALUES**

`audit()` returns:

- 0       on success.
- 1       on failure and sets `errno` to indicate the error.

**ERRORS**

- `EFAULT`       *record* points outside the process's allocated address space.
- `EINVAL`       The length specified in the audit record is too short, or more than `MAXAUDITDATA`.
- `EPERM`        The process's effective user ID is not super-user.

**SEE ALSO**

`auditsvc(2)`, `getaudit(2)`, `setuseraudit(2)`, `audit_args(3)`, `audit.log(5)`, `auditd(8)`

**NAME**

auditon – manipulate auditing

**SYNOPSIS**

```
#include <sys/label.h>
```

```
#include <sys/audit.h>
```

```
int auditon (condition)
```

```
int condition;
```

**DESCRIPTION**

The **auditon()** system call sets system auditing to the requested *condition* if and only if the current state of auditing allows that transition. Legitimate values for *condition* are:

AUC_UNSET	on/off has not been decided yet
AUC_AUDITING	auditing is to be done
AUC_NOAUDIT	auditing is not to be done

The permitted transitions are:

- Any condition may be changed back to itself.
- AUC\_UNSET may be changed to AUC\_AUDITING or AUC\_NOAUDIT.
- AUC\_AUDITING may be changed to AUC\_NOAUDIT.
- AUC\_NOAUDIT may be changed to AUC\_AUDITING.

Once changed, it is not possible to get back to AUC\_UNSET.

Only the super-user may successfully execute this call.

**RETURN VALUES**

**auditon()** returns the old audit condition value on success. On failure, it returns `-1` and sets **errno** to indicate the error.

**ERRORS**

EINVAL	The <i>condition</i> specified is outside the range of valid values.
	The current condition precludes the requested change.
EPERM	Neither of the process's effective or real user ID is super-user.

**SEE ALSO**

**audit(2)**, **setuseraudit(2)**

**NAME**

`auditsvc` – write audit records to specified file descriptor

**SYNOPSIS**

```
int auditsvc(fd, limit)
int fd;
int limit;
```

**DESCRIPTION**

The `auditsvc()` system call specifies the audit log file to the kernel. The kernel writes audit records to this file until an exceptional condition occurs and then the call returns. The parameter *fd* is a file descriptor that identifies the audit file. Programs should open this file for writing before calling `auditsvc()`. The parameter *limit* specifies a value between 0 and 100, instructing `auditsvc()` to return when the percentage of free disk space on the audit filesystem drops below this limit. Thus, the invoking program can take action to avoid running out of disk space. The `auditsvc()` system call does not return until one of the following conditions occurs:

- The process receives a signal that is not blocked or ignored.
- An error is encountered writing to the audit log file.
- The minimum free space (as specified by *limit*), has been reached.

Only processes with a real or effective user ID of super-user may execute this call successfully.

**RETURN VALUES**

`auditsvc()` returns only on an error.

**ERRORS**

EAGAIN	The descriptor referred to a <i>stream</i> , was marked for System V-style non-blocking I/O, and no data could be written immediately.
EBADF	<i>fd</i> is not a valid descriptor open for writing.
EBUSY	A second process attempted to perform this call. A second process attempted to perform this call.
EDQUOT	The user's quota of disk blocks on the file system containing the file has been exhausted. Audit filesystem space is below the specified limit.
EFBIG	An attempt was made to write a file that exceeds the process's file size limit or the maximum file size.
EINTR	The call is forced to terminate prematurely due to the arrival of a signal whose <code>SV_INTERRUPT</code> bit in <code>sv_flags</code> is set (see <code>sigvec(2)</code> ). <code>signal(3V)</code> , in the System V compatibility library, sets this bit for any signal it catches.
EINVAL	Auditing is disabled (see <code>auditon(2)</code> ). <i>fd</i> does not refer to a file of an appropriate type. Regular files are always appropriate.
EIO	An I/O error occurred while reading from or writing to the file system.
ENOSPC	There is no free space remaining on the file system containing the file.
ENXIO	A hangup occurred on the <i>stream</i> being written to.
EPERM	The process's effective or real user ID is not super-user.
EWouldBlock	The file was marked for 4.2BSD-style non-blocking I/O, and no data could be written immediately.

**SEE ALSO**

`audit(2)`, `auditon(2)`, `sigvec(2)`, `signal(3V)`, `audit.log(5)`, `auditd(8)`

**NAME**

**bind** – bind a name to a socket

**SYNOPSIS**

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

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

**DESCRIPTION**

**bind()** assigns a name to an unnamed socket. When a socket is created with **socket(2)** 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**

**bind()** returns:

- 0       on success.
- 1       on failure and sets **errno** to indicate the error.

**ERRORS**

- |               |  |
|---------------|--|
| 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 from the local machine.                                   |
| EBADF         | <i>s</i> is not a valid descriptor.  |
| EFAULT        | The <i>name</i> parameter is not in a valid part of the user address space.                      |
| EINVAL        | <i>namelen</i> is not the size of a valid address for the specified address family.              |
| ENOTSOCK      | <i>s</i> 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 path name in <i>name</i> .   |
| EIO          | An I/O error occurred while making the directory entry or allocating the inode.  |
| EISDIR       | A null path name was specified.  |
| ELOOP        | Too many symbolic links were encountered in translating the path name in <i>name</i> .   |
| ENAMETOOLONG | The length of the path argument exceeds {PATH_MAX}.<br>A pathname component is longer than {NAME_MAX} (see <b>sysconf(2V)</b> ) while {_POSIX_NO_TRUNC} is in effect (see <b>pathconf(2V)</b> ). |
| ENOENT       | A component of the path prefix of the path name in <i>name</i> does not exist.   |
| ENOTDIR      | A component of the path prefix of the path name in <i>name</i> is not a directory.   |
| EROFS        | The inode would reside on a read-only file system.   |

**SEE ALSO**

**connect(2), getsockname(2), listen(2), socket(2), unlink(2V)**



**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(2V)**),

The rules used in name binding vary between communication domains. Consult the manual entries in section 4 for detailed information.

**NAME**

**brk, sbrk** – change data segment size

**SYNOPSIS**

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

int brk(addr)
caddr_t addr;

caddr_t sbrk(incr)
int incr;
```

**DESCRIPTION**

**brk()** sets the system's idea of the lowest data segment location not used by the program (called the *break*) to *addr* (rounded up to the next multiple of the system's page size).

In the alternate function **sbrk()**, *incr* more bytes are added to the program's data space and a pointer to the start of the new area is returned.

When a program begins execution using **execve()** the break is set at the highest location defined by the program and data storage areas.

The **getrlimit(2)** system call may be used to determine the maximum permissible size of the *data* segment; it will not be possible to set the break beyond the **rlim\_max** value returned from a call to **getrlimit()**, that is to say, "**etext + rlim.rlim\_max**." (See **end(3)** for the definition of **etext()**.)

**RETURN VALUES**

**brk()** returns:

0        on success.

-1       on failure and sets **errno** to indicate the error.

**sbrk()** returns the old break value on success. On failure, it returns **(caddr\_t) -1** and sets **errno** to indicate the error.

**ERRORS**

**brk()** and **sbrk()** will fail and no additional memory will be allocated if one of the following occurs:

- |               |   |
|---------------|---|
| <b>ENOMEM</b> | The data segment size limit, as set by <b>setrlimit()</b> (see <b>getrlimit(2)</b> ), would be exceeded.<br>The maximum possible size of a data segment (compiled into the system) would be exceeded.<br>Insufficient space exists in the swap area to support the expansion.<br>Out of address space; the new break value would extend into an area of the address space defined by some previously established mapping (see <b>mmap(2)</b> ). |
|---------------|---|

**SEE ALSO**

**execve(2V)**, **mmap(2)**, **getrlimit(2)**, **malloc(3V)**, **end(3)**

**WARNINGS**

Programs combining the **brk()** and **sbrk()** system calls and **malloc()** will not work. Many library routines use **malloc()** internally, so use **brk()** and **sbrk()** only when you know that **malloc()** definitely will not be used by any library routine.

**BUGS**

Setting the break may fail due to a temporary lack of swap space. It is not possible to distinguish this from a failure caused by exceeding the maximum size of the data segment without consulting **getrlimit()**.

**NAME**

**chdir** – change current working directory

**SYNOPSIS**

```
int chdir(path)
char *path;

int fchdir(fd)
int fd;
```

**DESCRIPTION**

**chdir()** and **fchdir()** make the directory specified by *path* or *fd* the current working directory. Subsequent references to pathnames not starting with '/' are relative to the new current working directory.

In order for a directory to become the current directory, a process must have execute (search) access to the directory.

**RETURN VALUES**

**chdir()** returns:

0           on success.  
-1          on failure and sets **errno** to indicate the error.

**ERRORS**

<b>EACCES</b>	Search permission is denied for a component of the pathname.
<b>ENAMETOOLONG</b>	The length of the path argument exceeds {PATH_MAX}. A pathname component is longer than {NAME_MAX} while {_POSIX_NO_TRUNC} is in effect (see <b>pathconf(2V)</b> ).
<b>ENOENT</b>	The named directory does not exist.
<b>ENOTDIR</b>	A component of the pathname is not a directory.

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

**ENOENT**           *path* points to an empty string.

**WARNINGS**

**fchdir()** is provided as a performance enhancement and is guaranteed to fail under certain conditions. In particular, if auditing is active the call will never succeed, and **EINVAL** will be returned. Applications which use this system call must be coded to detect this failure and switch to using **chdir()** from that point on.

**NAME**

chmod, fchmod – change mode of file

**SYNOPSIS**

```
#include <sys/stat.h>

int chmod(path, mode)
char *path;
mode_t mode;

int fchmod(fd, mode)
int fd, mode;
```

**DESCRIPTION**

**chmod()** sets the mode of the file referred to by *path* or the descriptor *fd* according to *mode*. *mode* is the inclusive OR of the file mode bits (see **stat(2V)** for a description of these bits).

The effective user ID of the process must match the owner of the file or be super-user to change the mode of a file.

If the effective user ID of the process is not super-user and the process attempts to set the set group ID bit on a file owned by a group which is not in its supplementary group IDs, the **S\_ISGID** bit (set group ID on execution) is cleared.

If the **S\_ISVTX** (sticky) bit is set on a directory, an unprivileged user may not delete or rename files of other users in that directory.

If a user other than the super-user writes to a file, the set user ID and set group ID bits are turned off. This makes the system somewhat more secure by protecting set-user-ID (set-group-ID) files from remaining set-user-ID (set-group-ID) if they are modified, at the expense of a degree of compatibility.

**RETURN VALUES**

**chmod()** returns:

- 0       on success.
- 1       on failure and sets **errno** to indicate the error.

**ERRORS**

**chmod()** will fail and the file mode will be unchanged if:

- |              |  |
|--------------|--|
| EACCES       | Search permission is denied for a component of the path prefix of <i>path</i> .  |
| EFAULT       | <i>path</i> points outside the process's allocated address space.  |
| EINVAL       | <i>fd</i> refers to a socket, not to a file.   |
| EIO          | An I/O error occurred while reading from or writing to the file system.  |
| ELOOP        | Too many symbolic links were encountered in translating <i>path</i> .  |
| ENAMETOOLONG | The length of the path argument exceeds <b>{PATH_MAX}</b> .<br>A pathname component is longer than <b>{NAME_MAX}</b> while <b>{_POSIX_NO_TRUNC}</b> is in effect (see <b>pathconf(2V)</b> ). |
| ENOENT       | The file referred to by <i>path</i> does not exist.  |
| ENOTDIR      | A component of the path prefix of <i>path</i> is not a directory.  |
| EPERM        | The effective user ID does not match the owner of the file and the effective user ID is not the super-user.  |
| EROFS        | The file referred to by <i>path</i> resides on a read-only file system.  |

**fchmod()** will fail if:

- EBADF     The descriptor is not valid.

- EIO** An I/O error occurred while reading from or writing to the file system.
- EPERM** The effective user ID does not match the owner of the file and the effective user ID is not the super-user.
- EROFS** The file referred to by *fd* resides on a read-only file system.

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

- ENOENT** *path* points to a null pathname.

**SEE ALSO**

**chown(2V), open(2V), stat(2V), sticky(8)**

**BUGS**

**S\_ISVTX**, the “sticky bit”, is a misnomer, and is overloaded to mean different things for different file types.

**NAME**

chown, fchown – change owner and group of a file

**SYNOPSIS**

```
int chown(path, owner, group)
```

```
char *path;
```

```
int owner;
```

```
int group;
```

```
int fchown(fd, owner, group)
```

```
int fd;
```

```
int owner;
```

```
int group;
```

**SYSTEM V SYNOPSIS**

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

```
int chown(path, owner, group)
```

```
char *path;
```

```
uid_t owner;
```

```
gid_t group;
```

**DESCRIPTION**

The file that is named by *path* or referenced by *fd* has its *owner* and *group* changed as specified. Only the super-user may change the owner of the file, because if users were able to give files away, they could defeat the file-space accounting procedures (see NOTES). The owner of the file may change the group to a group of which he is a member. The super-user may change the group arbitrarily.

**fchown()** is particularly useful when used in conjunction with the file locking primitives (see **flock(2)**).

If *owner* or *group* is specified as  $-1$ , the corresponding ID of the file is not changed.

If a process whose effective user ID is not super-user successfully changes the group ID of a file, the set-user-ID and set-group-ID bits of the file mode, **S\_ISUID** and **S\_ISGID** respectively (see **stat(2V)**), will be cleared.

If the final component of *path* is a symbolic link, the ownership and group of the symbolic link is changed, not the ownership and group of the file or directory to which it points.

**RETURN VALUES**

**chown()** and **fchown()** return:

0        on success.

$-1$      on failure and set **errno** to indicate the error.

**ERRORS**

**chown()** will fail and the file will be unchanged if:

**EACCES**            Search permission is denied for a component of the path prefix of *path*.

**EFAULT**            *path* points outside the process's allocated address space.

**EIO**                An I/O error occurred while reading from or writing to the file system.

**ELOOP**            Too many symbolic links were encountered in translating *path*.

**ENAMETOOLONG**     The length of the path argument exceeds **{PATH\_MAX}**.

A pathname component is longer than **{NAME\_MAX}** (see **sysconf(2V)**) while **{\_POSIX\_NO\_TRUNC}** is in effect (see **pathconf(2V)**).

**ENOENT**            The file referred to by *path* does not exist.

**ENOTDIR**           A component of the path prefix of *path* is not a directory.

EPERM	The user ID specified by <i>owner</i> is not the current owner ID of the file. The group ID specified by <i>group</i> is not the current group ID of the file and is not in the process' supplementary group IDs, and the effective user ID is not the super-user.
EROFS	The file referred to by <i>path</i> resides on a read-only file system.
<b>fchown()</b> will fail if:	
EBADF	<i>fd</i> does not refer to a valid descriptor.
EINVAL	<i>fd</i> refers to a socket, not a file.
EIO	An I/O error occurred while reading from or writing to the file system.
EPERM	The user ID specified by <i>owner</i> is not the current owner ID of the file. The group ID specified by <i>group</i> is not the current group ID of the file and is not in the supplementary group IDs, and the effective user ID is not the super-user.
EROFS	The file referred to by <i>fd</i> resides on a read-only file system.

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

ENOENT                *path* points to an empty string.

**SEE ALSO**

**chmod(2V), flock(2)**

**NOTES**

For **chown()** to behave as described above, `{_POSIX_CHOWN_RESTRICTED}` must be in effect (see **pathconf(2V)**). `{_POSIX_CHOWN_RESTRICTED}` is always in effect on SunOS systems, but for portability, applications should call **pathconf()** to determine whether `{_POSIX_CHOWN_RESTRICTED}` is in effect for *path*.

If `{_POSIX_CHOWN_RESTRICTED}` is in effect for the file system on which the file referred to by *path* or *fd* resides, only the super-user may change the owner of the file. Otherwise, processes with effective user ID equal to the file owner or super-user may change the owner of the file.

**NAME**

chroot – change root directory

**SYNOPSIS**

```
int chroot(dirname)
char *dirname;

int fchroot(fd)
int fd;
```

**DESCRIPTION**

**chroot()** and **fchroot()** cause a directory to become the root directory, the starting point for path names beginning with '/'. The current working directory is unaffected by this call. This root directory setting is inherited across **execve(2V)** and by all children of this process created with **fork(2V)** calls.

In order for a directory to become the root directory a process must have execute (search) access to the directory and either the effective user ID of the process must be super-user or the target directory must be the system root or a loop-back mount of the system root (see **lofs(4S)**). **fchroot()** is further restricted in that while it is always possible to change to the system root using this call, it is not guaranteed to succeed in any other case, even should *fd* be in all respects valid.

The *dirname* argument to **chroot()** points to a path name of a directory. The *fd* argument to **fchroot()** is the open file descriptor of the directory which is to become the root.

The .. entry in the root directory is interpreted to mean the root directory itself. Thus, .. cannot be used to access files outside the subtree rooted at the root directory. Instead, **fchroot()** can be used to set the root back to a directory which was opened before the root directory was changed.

**WARNINGS**

The only use of **fchroot()** that is appropriate is to change back to the system root. While it may succeed in some other cases, it is guaranteed to fail if auditing is enabled. Super-user processes are not exempt from this limitation.

**RETURN VALUES**

**chroot()** returns:

- 0 on success.
- 1 on failure and sets **errno** to indicate the error.

**ERRORS**

**chroot()** will fail and the root directory will be unchanged if one or more of the following are true:

EACCES	Search permission is denied for a component of the path prefix of <i>dirname</i> . Search permission is denied for the directory referred to by <i>dirname</i> .
EBADF	The descriptor is not valid.
EFAULT	<i>dirname</i> points outside the process's allocated address space.
EINVAL	<b>fchroot()</b> attempted to change to a directory which is not the system root and external circumstances, such as auditing, do not allow this.
EIO	An I/O error occurred while reading from or writing to the file system.
ELOOP	Too many symbolic links were encountered in translating <i>dirname</i> .
ENAMETOOLONG	The length of the path argument exceeds {PATH_MAX}. A pathname component is longer than {NAME_MAX} (see <b>sysconf(2V)</b> ) while {_POSIX_NO_TRUNC} is in effect (see <b>pathconf(2V)</b> ).
ENOENT	The directory referred to by <i>dirname</i> does not exist.



ENOTDIR

A component of the path prefix of *dirname* is not a directory.  
The file referred to by *dirname* is not a directory.

EPERM

The effective user ID is not super-user.

**SEE ALSO**

**chdir(2V), execve(2V), fork(2V), lofs(4S)**

**NAME**

close – delete a descriptor

**SYNOPSIS**

```
int close (fd)
int fd;
```

**DESCRIPTION**

**close()** deletes a descriptor from the per-process object reference table. If *fd* is the last reference to the underlying object, then the object will be deactivated. For example, on the last close of a file the current *seek* pointer associated with the file is lost. On the last close of a socket (see **socket(2)**), associated naming information and queued data are discarded. On the last close of a file holding an advisory lock applied by **flock(2)**, the lock is released. (Record locks applied to the file by **lockf(3)**, however, are released on *any* call to **close()** regardless of whether *fd* is the last reference to the underlying object.)

**close()** does not unmap any mapped pages of the object referred to by *fd* (see **mmap()**, **munmap(2)**).

A close of all of a process's descriptors is automatic on **exit()**, but since there is a limit on the number of active descriptors per process, **close()** is necessary for programs that deal with many descriptors.

When a process forks (see **fork(2v)**), all descriptors for the new child process reference the same objects as they did in the parent before the fork. If a new process is then to be run using **execve(2V)**, the process would normally inherit these descriptors. Most of the descriptors can be rearranged with **dup(2V)** or deleted with **close()** before the **execve()** is attempted, but if some of these descriptors will still be needed if the **execve()** fails, it is necessary to arrange for them to be closed if the **execve()** succeeds. The **fcntl(2V)** operation **F\_SETFD** can be used to arrange that a descriptor will be closed after a successful **execve()**, or to restore the default behavior, which is to not close the descriptor.

If a STREAMS (see **intro(2)**) file is closed, and the calling process had previously registered to receive a **SIGPOLL** signal (see **sigvec(2)**) for events associated with that file (see **I\_SETSIG** in **streamio(4)**), the calling process will be unregistered for events associated with the file. The last **close()** for a stream causes that stream to be dismantled. If the descriptor is not marked for no-delay mode and there have been no signals posted for the stream, **close()** waits up to 15 seconds, for each module and driver, for any output to drain before dismantling the stream. If the descriptor is marked for no-delay mode or if there are any pending signals, **close()** does not wait for output to drain, and dismantles the stream immediately.

**RETURN VALUES**

**close()** returns:

- 0       on success.
- 1       on failure and sets **errno** to indicate the error.

**ERRORS**

- EBADF**       *fd* is not an active descriptor.
- EINTR**       A signal was caught before the close completed.

**SEE ALSO**

**accept(2)**, **dup(2V)**, **execve(2V)**, **fcntl(2V)**, **flock(2)**, **intro(2)**, **open(2V)**, **pipe(2V)**, **sigvec(2)**, **socket(2)**, **socketpair(2)**, **streamio(4)**

**NAME**

connect – initiate a connection on a socket

**SYNOPSIS**

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

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

**DESCRIPTION**

The parameter *s* is a socket. If it is of type `SOCK_DGRAM`, then this call specifies the peer with which the socket is to be associated; this address is that to which datagrams are to be sent, and the only address from which datagrams are to be received. If it is of type `SOCK_STREAM`, then this call attempts to make a connection to another socket. The other socket is specified by *name* which is an address in the communications space of the socket. Each communications space interprets the *name* parameter in its own way. Generally, stream sockets may successfully `connect()` only once; datagram sockets may use `connect()` multiple times to change their association. Datagram sockets may dissolve the association by connecting to an invalid address, such as a null address.

**RETURN VALUES**

`connect()` returns:

- 0        on success.
- 1       on failure and sets `errno` to indicate the error.

**ERRORS**

The call fails if:

<code>EADDRINUSE</code>	The address is already in use.
<code>EADDRNOTAVAIL</code>	The specified address is not available on the remote machine.
<code>EAFNOSUPPORT</code>	Addresses in the specified address family cannot be used with this socket.
<code>EALREADY</code>	The socket is non-blocking and a previous connection attempt has not yet been completed.
<code>EBADF</code>	<i>s</i> is not a valid descriptor.
<code>ECONNREFUSED</code>	The attempt to connect was forcefully rejected. The calling program should <code>close(2V)</code> the socket descriptor, and issue another <code>socket(2)</code> call to obtain a new descriptor before attempting another <code>connect(2)</code> call.
<code>EFAULT</code>	The <i>name</i> parameter specifies an area outside the process address space.
<code>EINPROGRESS</code>	The socket is non-blocking and the connection cannot be completed immediately. It is possible to <code>select(2)</code> for completion by selecting the socket for writing.
<code>EINTR</code>	The connection attempt was interrupted before any data arrived by the delivery of a signal.
<code>EINVAL</code>	<i>namelen</i> is not the size of a valid address for the specified address family.
<code>EISCONN</code>	The socket is already connected.
<code>ENETUNREACH</code>	The network is not reachable from this host.
<code>ENOTSOCK</code>	<i>s</i> is a descriptor for a file, not a socket.
<code>ETIMEDOUT</code>	Connection establishment timed out without establishing a connection.

The following errors are specific to connecting names in the UNIX domain. These errors may not apply in future versions of the UNIX IPC domain.

EACCES	Search permission is denied for a component of the path prefix of the path name in <i>name</i> .
ELOOP	Too many symbolic links were encountered in translating the path name in <i>name</i> .
EIO	An I/O error occurred while reading from or writing to the file system.
ENAMETOOLONG	The length of the path argument exceeds {PATH_MAX}. A pathname component is longer than {NAME_MAX} (see <code>sysconf(2V)</code> ) while {_POSIX_NO_TRUNC} is in effect (see <code>pathconf(2V)</code> ).
ENOENT	A component of the path prefix of the path name in <i>name</i> does not exist. The socket referred to by the path name in <i>name</i> does not exist.
ENOTDIR	A component of the path prefix of the path name in <i>name</i> is not a directory.
ENOTSOCK	The file referred to by <i>name</i> is not a socket.
EPROTOTYPE	The file referred to by <i>name</i> is a socket of a type other than the type of <i>s</i> (e.g., <i>s</i> is a SOCK_DGRAM socket, while <i>name</i> refers to a SOCK_STREAM socket).

SEE ALSO

`accept(2)`, `close(2V)`, `connect(2)`, `getsockname(2)`, `select(2)`, `socket(2)`

**NAME**

**creat** – create a new file

**SYNOPSIS**

```
int creat(path, mode)
char *path;
int mode;
```

**SYSTEM V SYNOPSIS**

```
#include <sys/stat.h>

int creat(path, mode)
char *path;
mode_t mode;
```

**DESCRIPTION**

This interface is made obsolete by **open(2V)**, since,

```
creat(path, mode);
```

is equivalent to

```
open(path, O_WRONLY | O_CREAT | O_TRUNC, mode);
```

**creat()** creates a new ordinary file or prepares to rewrite an existing file named by the pathname pointed to by *path*. If the file did not exist, it is given the mode *mode*, as modified by the process's mode mask (see **umask(2V)**). See **stat(2V)** for the construction of *mode*.

If the file exists, its mode and owner remain unchanged, but it is truncated to 0 length. Otherwise, the file's owner ID is set to the effective user ID of the process, and upon successful completion, **creat()** marks for update the **st\_atime**, **st\_ctime**, and **st\_mtime** fields of the file (see **stat(2V)**) and the **st\_ctime** and **st\_mtime** fields of the parent directory.

The file's group ID is set to either:

- the effective group ID of the process, if the filesystem was not mounted with the BSD file-creation semantics flag (see **mount(2V)**) and the set-gid bit of the parent directory is clear, or
- the group ID of the directory in which the file is created.

The low-order 12 bits of the file mode are set to the value of *mode*, modified as follows:

- All bits set in the process's file mode creation mask are cleared. See **umask(2V)**.
- The "save text image after execution" (sticky) bit of the mode is cleared. See **chmod(2V)**.
- The "set group ID on execution" bit of the mode is cleared if the effective user ID of the process is not super-user and the process is not a member of the group of the created file.

Upon successful completion, the file descriptor is returned and the file is open for writing, even if the access permissions of the file mode do not permit writing. The file pointer is set to the beginning of the file. The file descriptor is set to remain open across **execve(2V)** system calls. See **fcntl(2V)**.

If the file did not previously exist, upon successful completion, **creat()** marks for update the **st\_ctime** and **st\_mtime** fields of the file and the **st\_ctime** and **st\_mtime** fields of the parent directory.

**RETURN VALUES**

**creat()** returns a non-negative descriptor that only permits writing on success. On failure, it returns **-1** and sets **errno** to indicate the error.

**ERRORS**

**EACCES**

Search permission is denied for a component of the path prefix.

The file referred to by *path* does not exist and the directory in which it is to be created is not writable.

The file referred to by *path* exists, but it is unwritable.

EDQUOT	The directory in which the entry for the new file is being placed cannot be extended because the user's quota of disk blocks on the file system containing the directory has been exhausted. The user's quota of inodes on the file system on which the file is being created has been exhausted.
EFAULT	<i>path</i> points outside the process's allocated address space.
EINTR	The <code>creat()</code> operation was interrupted by a signal.
EIO	An I/O error occurred while making the directory entry or allocating the inode.
EISDIR	The file referred to by <i>path</i> is a directory.
ELOOP	Too many symbolic links were encountered in translating the pathname pointed to by <i>path</i> .
EMFILE	There are already too many files open.
ENAMETOOLONG	The length of the path argument exceeds <code>{PATH_MAX}</code> . A pathname component is longer than <code>{NAME_MAX}</code> while <code>{_POSIX_NO_TRUNC}</code> is in effect (see <code>pathconf(2V)</code> ).
ENFILE	The system file table is full.
ENOENT	A component of the path prefix does not exist.
ENOSPC	The directory in which the entry for the new file is being placed cannot be extended because there is no space left on the file system containing the directory. There are no free inodes on the file system on which the file is being created.
ENOTDIR	A component of the path prefix is not a directory.
ENXIO	The file is a character special or block special file, and the associated device does not exist.
EOPNOTSUPP	The file was a socket (not currently implemented).
EROFS	The file referred to by <i>path</i> resides, or would reside, on a read-only file system.

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

ENOENT            *path* points to an empty string.

**SEE ALSO**

`close(2V)`, `chmod(2V)`, `execve(2V)`, `fcntl(2V)`, `flock(2)`, `mount(2V)`, `open(2V)`, `write(2V)`, `umask(2V)`

**NOTES**

The *mode* given is arbitrary; it need not allow writing. This feature has been used in the past by programs to construct a simple exclusive locking mechanism. It is replaced by the `O_EXCL` open mode, or `flock(2)` facility.

**NAME**

**dup, dup2** – duplicate a descriptor

**SYNOPSIS**

```
int dup(fd)  
int fd;
```

```
int dup2(fd1, fd2)  
int fd1, fd2;
```

**DESCRIPTION**

**dup()** duplicates an existing object descriptor. The argument *fd* is a small non-negative integer index in the per-process descriptor table. The value must be less than the size of the table, which is returned by **getdtablesize(2)**. The new descriptor returned by the call is the lowest numbered descriptor that is not currently in use by the process.

With **dup2()**, *fd2* specifies the desired value of the new descriptor. If descriptor *fd2* is already in use, it is first deallocated as if it were closed by **close(2V)**.

The new descriptor has the following in common with the original:

- It refers to the same object that the old descriptor referred to.
- It uses the same seek pointer as the old descriptor. (that is, both file descriptors share one seek pointer).
- It has the same access mode (read, write or read/write) as the old descriptor.

Thus if *fd2* and *fd1* are duplicate references to an open file, **read(2V)**, **write(2V)**, and **lseek(2V)** calls all move a single seek pointer into the file, and append mode, non-blocking I/O and asynchronous I/O options are shared between the references. If a separate seek pointer into the file is desired, a different object reference to the file must be obtained by issuing an additional **open(2V)** call. The close-on-exec flag on the new file descriptor is unset.

The new file descriptor is set to remain open across **exec** system calls (see **fcntl(2V)**).

**RETURN VALUES**

**dup()** and **dup2()** return a new descriptor on success. On failure, they return **-1** and set **errno** to indicate the error.

**ERRORS**

**EBADF**            *fd1* or *fd2* is not a valid active descriptor.  
**EMFILE**          Too many descriptors are active.

**SEE ALSO**

**accept(2)**, **close(2V)**, **fcntl(2V)**, **getdtablesize(2)**, **lseek(2V)**, **open(2V)**, **pipe(2V)**, **read(2V)**, **socket(2)**, **socketpair(2)**, **write(2V)**

## NAME

`execve` – execute a file

## SYNOPSIS

```
int execve(path, argv, envp)
char *path, *argv[], *envp[];
```

## DESCRIPTION

`execve()` transforms the calling process into a new process. The new process is constructed from an ordinary file, whose name is pointed to by *path*, called the *new process file*. This file is either an executable object file, or a file of data for an interpreter. An executable object file consists of an identifying header, followed by pages of data representing the initial program (text) and initialized data pages. Additional pages may be specified by the header to be initialized with zero data. See `a.out(5)`.

An interpreter file begins with a line of the form `#! interpreter [arg]`. Only the first thirty-two characters of this line are significant. When *path* refers to an interpreter file, `execve()` invokes the specified *interpreter*. If the optional *arg* is specified, it becomes the first argument to the *interpreter*, and the pathname to which *path* points becomes the second argument. Otherwise, the pathname to which *path* points becomes the first argument. The original arguments are shifted over to become the subsequent arguments. The zeroth argument, normally the pathname to which *path* points, is left unchanged.

There can be no return from a successful `execve()` because the calling process image is lost. This is the mechanism whereby different process images become active.

The argument *argv* is a pointer to a null-terminated array of character pointers to null-terminated character strings. These strings constitute the argument list to be made available to the new process. By convention, at least one argument must be present in this array, and the first element of this array should be the name of the executed program (that is, the last component of *path*).

The argument *envp* is also a pointer to a null-terminated array of character pointers to null-terminated strings. These strings pass information to the new process which are not directly arguments to the command (see `environ(5V)`).

The number of bytes available for the new process's combined argument and environment lists (including null terminators, pointers and alignment bytes) is `{ARG_MAX}` (see `sysconf(2V)`). On SunOS systems, `{ARG_MAX}` is currently one megabyte.

Descriptors open in the calling process remain open in the new process, except for those for which the close-on-exec flag is set (see `close(2V)` and `fcntl(2V)`). Descriptors which remain open are unaffected by `execve()`.

Signals set to the default action (`SIG_DFL`) in the calling process image are set to the default action in the new process image. Signals set to be ignored (`SIG_IGN`) by the calling process image are ignored by the new process image. Signals set to be caught by the calling process image are reset to the default action in the new process image. Signals set to be blocked in the calling process image remain blocked in the new process image, regardless of changes to the signal action. The signal stack is reset to be undefined (see `sigvec(2)` for more information).

Each process has a *real* user ID and group ID and an *effective* user ID and group ID. The *real* ID identifies the person using the system; the *effective* ID determines their access privileges. `execve()` changes the effective user or group ID to the owner or group of the executed file if the file has the "set-user-ID" or "set-group-ID" modes. The *real* UID and GID are not affected. The effective user ID and effective group ID of the new process image are saved as the saved set-user-ID and saved set-group-ID respectively, for use by `setuid(3V)`.

`execve()` sets the `SEXECED` flag for the new process image (see `setpgid(2V)`).

The shared memory segments attached to the calling process will not be attached to the new process (see `shmop(2)`).



Profiling is disabled for the new process; see **profil(2)**.

Upon successful completion, **execve()** marks for update the **st\_atime** field of the file. **execve()** also marks **st\_atime** for update if it fails, but is able find the process image file.

If **execve()** succeeds, the process image file is considered to have been opened (see **open(2V)**). The corresponding close (see **close(2V)**) is considered to occur after the open, but before process termination or successful completion of a subsequent call to **execve()**.

The new process also inherits the following attributes from the calling process:

<i>attribute</i>	<i>see</i>
process ID	<b>getpid(2)</b>
parent process ID	<b>getpid(2)</b>
process group ID	<b>getpgrp(2V)</b> , <b>setpgid(2V)</b>
session membership	<b>setsid(2)</b>
real user ID	<b>getuid(2)</b>
real group ID	<b>getgid(2)</b>
supplementary group IDs	<b>Intro(2)</b>
time left until an alarm	<b>alarm(3C)</b>
supplementary group IDs	<b>getgroups(2)</b>
semadj values	<b>semop(2)</b>
working directory	<b>chdir(2)</b>
root directory	<b>chroot(2)</b>
controlling terminal	<b>termio(4)</b>
trace flag	<b>ptrace(2)</b> , request 0
resource usages	<b>getrusage(2)</b>
interval timers	<b>getitimer(2)</b>
resource limits	<b>getrlimit(2)</b>
file mode mask	<b>umask(2)</b>
process signal mask	<b>sigvec(2)</b> , <b>sigprocmask(2V)</b> , <b>sigsetmask(2)</b>
pending signals	<b>sigpending(2)</b>
<b>tms_utime</b> , <b>tms_stime</b> , <b>tms_cutime</b> , <b>tms_cstime</b>	<b>times(3C)</b>

When the executed program begins, it is called as follows:

```
main(argc, argv, envp)
int argc;
char *argv[ ], *envp[ ];
```

where *argc* is the number of elements in *argv* (the “arg count”, not counting the NULL terminating pointer) and *argv* points to the array of character pointers to the arguments themselves.

*envp* is a pointer to an array of strings that constitute the *environment* of the process. A pointer to this array is also stored in the global variable **environ**. Each string consists of a name, an “=”, and a null-terminated value. The array of pointers is terminated by a NULL pointer. The shell **sh(1)** passes an environment entry for each global shell variable defined when the program is called. See **environ(5V)** for some conventionally used names.

Note: Passing values for *argc*, *argv*, and *envp* to **main()** is optional.

#### RETURN VALUES

**execve()** returns to the calling process only on failure. It returns **-1** and sets **errno** to indicate the error.

#### ERRORS

- |        |  |
|--------|--|
| E2BIG  | The total number of bytes in the new process file’s argument and environment lists exceeds { <b>ARG_MAX</b> } (see <b>sysconf(2V)</b> ). |
| EACCES | Search permission is denied for a component of the new process file’s path prefix.   |

	The new process file is not a regular file.
	Execute permission is denied for the new process file.
EFAULT	The new process file is not as long as indicated by the size values in its header. <i>path</i> , <i>argv</i> , or <i>envp</i> points to an illegal address.
EIO	An I/O error occurred while reading from the file system.
ENAMETOOLONG	The length of the path argument exceeds {PATH_MAX}. A pathname component is longer than {NAME_MAX} (see <code>sysconf(2V)</code> ) while {_POSIX_NO_TRUNC} is in effect (see <code>pathconf(2V)</code> ).
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
ENOENT	One or more components of the path prefix of the new process file does not exist. The new process file does not exist.
ENOEXEC	The new process file has the appropriate access permission, but has an invalid magic number in its header.
ENOMEM	The new process file requires more virtual memory than is allowed by the imposed maximum ( <code>getrlimit(2)</code> ).
ENOTDIR	A component of the path prefix of the new process file is not a directory.

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

ENOENT            *path* points to a null pathname.

**SEE ALSO**

`sh(1)`, `chdir(2V)`, `chroot(2)`, `close(2V)`, `exit(2V)`, `fcntl(2V)`, `fork(2V)`, `getgroups(2V)`, `getitimer(2)`, `getpid(2V)`, `getrlimit(2)`, `getrusage(2)`, `profil(2)`, `ptrace(2)`, `semop(2)`, `getpgrp(2V)`, `shmop(2)`, `sigvec(2)`, `execl(3V)`, `setuid(3V)`, `termio(4)`, `a.out(5)`, `environ(5V)`

**WARNINGS**

If a program is `setuid()` to a non-super-user, but is executed when the real user ID is super-user, then the program has some of the powers of a super-user as well.

**NAME**

**\_exit** – terminate a process

**SYNOPSIS**

```
void _exit(status)
int status;
```

**DESCRIPTION**

**\_exit()** terminates a process with the following consequences:

All of the descriptors open in the calling process are closed. This may entail delays, for example, waiting for output to drain; a process in this state may not be killed, as it is already dying.

If the parent process of the calling process is executing a **wait()** or **waitpid()**, or is interested in the **SIGCHLD** signal, then it is notified of the calling process's termination and the low-order eight bits of *status* are made available to it (see **wait(2V)**).

If the parent process of the calling process is not executing a **wait()** or **waitpid()**, *status* is saved for return to the parent process whenever the parent process executes an appropriate subsequent **wait()** or **waitpid()**.

The parent process ID of all of the calling process's existing child processes are also set to 1. This means that the initialization process (see **intro(2)**) inherits each of these processes as well. Any stopped children are restarted with a hangup signal (**SIGHUP**).

If the process is a controlling process, **SIGHUP** is sent to each process in the foreground process group of the controlling terminal belonging to the calling process, and the controlling terminal associated with the session is disassociated from the session, allowing it to be acquired by a new controlling process (see **setsid(2V)**).

If **\_exit()** causes a process group to become orphaned, and if any member of the newly-orphaned process group is stopped, then **SIGHUP** followed by **SIGCONT** is sent to each process in the newly-orphaned process group (see **setpgid(2V)**).

Each attached shared memory segment is detached and the value of **shm\_nattach** in the data structure associated with its shared memory identifier is decremented by 1.

For each semaphore for which the calling process has set a *semadj* value (see **semop(2)**), that *semadj* value is added to the *semval* of the specified semaphore.

If process accounting is enabled (see **acct(2V)**), an accounting record is written to the accounting file.

Most C programs will call the library routine **exit(3)** which performs cleanup actions in the standard I/O library before calling **\_exit()**.

**RETURN VALUES**

**\_exit()** never returns.

**SEE ALSO**

**intro(2)**, **acct(2V)**, **fork(2V)**, **semop(2)**, **wait(2V)**, **exit(3)**

## NAME

fcntl – file control

## SYNOPSIS

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

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

```
#include <fcntl.h>
```

```
int fcntl(fd, cmd, arg)
```

```
int fd, cmd, arg;
```

## DESCRIPTION

**fcntl()** performs a variety of functions on open descriptors. The argument *fd* is an open descriptor used by *cmd* as follows:

**F\_DUPFD** Returns a new descriptor, which has the smallest value greater than or equal to *arg*. It refers to the same object as the original descriptor, and has the same access mode (read, write or read/write). The new descriptor shares descriptor status flags with *fd*, and if the object was a file, the same file pointer. It is also associated with a **FD\_CLOEXEC** (close-on-exec) flag set to remain open across **execve(2V)** system calls.

**F\_GETFD** Get the **FD\_CLOEXEC** (close-on-exec) flag associated with *fd*. If the low-order bit is 0, the file remains open after executing **execve()**, otherwise it is closed.

**F\_SETFD** Set the **FD\_CLOEXEC** (close-on-exec) flag associated with *fd* to the low order bit of *arg* (0 or 1 as above).

Note: this is a per-process and per-descriptor flag. Setting or clearing it for a particular descriptor does not affect the flag on descriptors copied from it by **dup(2V)** or **F\_DUPFD**, nor does it affect the flag on other processes of that descriptor.

**F\_GETFL** Get descriptor status flags (see **fcntl(5)** for definitions).

**F\_SETFL** Set descriptor status flags (see **fcntl(5)** for definitions). The following flags are the only ones whose values may change: **O\_APPEND**, **O\_SYNC**, and **O\_NDELAY**, and the **FASYNC**, **FNDELAY**, and **FNDELAY** flags defined in **<fcntl.h>**.

**O\_NDELAY** and **FNDELAY** are identical.

Descriptor status flag values set by **F\_SETFL** affects descriptors copied using **dup(2V)**, **F\_DUPFD** or other processes.

Setting or clearing the **FNDELAY** flag on a descriptor causes an **FIONBIO ioctl(2)** request to be performed on the object referred to by that descriptor. Setting or clearing non-blocking mode, and setting or clearing the **FASYNC** flag on a descriptor causes an **FIOASYNC ioctl(2)** request to be performed on the object referred to by that descriptor, setting or clearing asynchronous mode. Thus, all descriptors referring to the object are affected.

**F\_GETLK** Get a description of the first lock which would block the lock specified by the **flock** structure pointed to by *arg* (see the definition of **struct flock** below). If a lock exists, The **flock** structure is overwritten with that lock's description. Otherwise, the structure is passed back with the lock type set to **F\_UNLOCK** and is otherwise unchanged.

**F\_SETLK** Set or clear a file segment lock according to the **flock** structure pointed to by *arg*. **F\_SETLK** is used to set shared (**F\_RDLCK**) or exclusive (**F\_WRLCK**) locks, or to remove those locks (**F\_UNLCK**). If the specified lock cannot be applied, **fcntl()** fails and returns immediately.

- F\_SETLKW** This *cmd* is the same as **F\_SETLK** except that if a shared or exclusive lock is blocked by other locks, the process waits until the requested lock can be applied. If a signal that is set to be caught (see **signal(3V)**) is received while **fcntl()** is waiting for a region, the call to **fcntl()** is interrupted. Upon return from the process's signal handler, **fcntl()** fails and returns, and the requested lock is not applied.
- F\_GETOWN** Get the process ID or process group currently receiving **SIGIO** and **SIGURG** signals; process groups are returned as negative values.
- F\_SETOWN** Set the process or process group to receive **SIGIO** and **SIGURG** signals. Process groups are specified by supplying *arg* as negative, otherwise *arg* is interpreted as a process ID.
- F\_RSETLK**  
**F\_RSETLKW**  
**F\_RGETLK** Are used by the network lock daemon, **lockd(8C)**, to communicate with the NFS server kernel to handle locks on the NFS files.

Record locking is done with either *shared* (**F\_RDLCK**), or *exclusive* (**F\_WRLCK**) locks. More than one process may hold a shared lock on a particular file segment, but if one process holds an exclusive lock on the segment, no other process may hold any lock on the segment until the exclusive lock is removed.

In order to claim a shared lock, a descriptor must be opened with read access. Descriptors for exclusive locks must be opened with write access.

A shared lock may be changed to an exclusive lock, and vice versa, simply by specifying the appropriate lock type with a **F\_SETLK** or **F\_SETLKW** *cmd*. Before the previous lock is released and the new lock applied, any other processes already in line must gain and release their locks.

If *cmd* is **F\_SETLKW** and the requested lock cannot be claimed immediately (for instance, when another process holds an exclusive lock that overlaps the current request) the calling process is blocked until the lock may be acquired. These blocks may be interrupted by signals. Care should be taken to avoid deadlocks caused by multiple processes all blocking the same records.

A shared or exclusive lock is either *advisory* or *mandatory* depending on the mode bits of the file containing the locked segment. The lock is mandatory if the set-GID bit (**S\_ISGID**) is set and the group execute bit (**S\_IXGRP**) is clear (see **stat(2V)** for information about mode bits). Otherwise, the lock is advisory.

If a process holds a mandatory shared lock on a segment of a file, other processes may read from the segment, but write operations block until all locks are removed. If a process holds a mandatory exclusive lock on a segment of a file, both read and write operations block until the lock is removed (see **WARNINGS**).

An advisory lock does not affect read and write access to the locked segment. Advisory locks may be used by cooperating processes checking for locks using **F\_GETLCK** and voluntarily observing the indicated read and write restrictions.

The record to be locked or unlocked is described by the **flock** structure defined in **<fcntl.h>** as follows:

```

struct flock {
    short l_type;    /* F_RDLCK, F_WRLCK, or F_UNLCK */
    short l_whence; /* flag to choose starting offset */
    long  l_start;   /* relative offset, in bytes */
    long  l_len;     /* length, in bytes; 0 means lock to EOF */
    pid_t l_pid;     /* returned with F_GETLK */
};

```

The **flock** structure describes the type (**l\_type**), starting offset (**l\_whence**), relative offset (**l\_start**), and size (**l\_len**) of the file segment to be affected. **l\_whence** is set to **SEEK\_SET**, **SEEK\_CUR**, or **SEEK\_END** (see **lseek(2V)**) to indicate that the relative offset is to be measured from the start of the file, current position, or EOF, respectively. The process id field (**l\_pid**) is only used with the **F\_GETLK** *cmd* to return the description of a lock held by another process. Note: do not confuse **struct flock** with the function **flock(2)**. They are unrelated.

Locks may start or extend beyond the current EOF, but may not be negative relative to the beginning of the file. Setting `l_len` to zero (0) extends the lock to EOF. If `l_whence` is set to `SEEK_SET` and `l_start` and `l_len` are set to zero (0), the entire file is locked. Changing or unlocking the subset of a locked segment leaves the smaller segments at either end locked. Locking a segment already locked by the calling process causes the old lock type to be removed and the new lock type to take affect. All locks associated with a file for a given process are removed when the file is closed or the process terminates. Locks are not inherited by the child process in a `fork(2V)` system call.

`fcntl()` record locks are implemented in the kernel for local locks, and throughout the network by the network lock daemon (`lockd(8C)`) for remote locks on NFS files. If the file server crashes and has to be rebooted, the lock daemon attempts to recover all locks that were associated with that server. If a lock cannot be reclaimed, the process that held the lock is issued a `SIGLOST` signal.

In order to maintain consistency in the network case, data must not be cached on client machines. For this reason, file buffering for an NFS file is turned off when the first lock is attempted on the file. Buffering remains off as long as the file is open. Programs that do I/O buffering in the user address space, however, may have inconsistent results. The standard I/O package, for instance, is a common source of unexpected buffering.

#### SYSTEM V DESCRIPTION

`O_NDELAY` and `FNBIO` are identical.

#### RETURN VALUES

On success, the value returned by `fcntl()` depends on `cmd` as follows:

<code>F_DUPFD</code>	A new descriptor.
<code>F_GETFD</code>	Value of flag (only the low-order bit is defined).
<code>F_GETFL</code>	Value of flags.
<code>F_GETOWN</code>	Value of descriptor owner.
other	Value other than -1.

On failure, `fcntl()` returns -1 and sets `errno` to indicate the error.

#### ERRORS

<code>EACCES</code>	<code>cmd</code> is <code>F_SETLK</code> , the lock type ( <code>l_type</code> ) is <code>F_RDLCK</code> (shared lock), and the file segment to be locked is already under an exclusive lock held by another process. This error is also returned if the lock type is <code>F_WRLCK</code> (exclusive lock) and the file segment is already locked with a shared or exclusive lock.  Note: In future, <code>fcntl()</code> may generate <code>EAGAIN</code> under these conditions, so applications testing for <code>EACCES</code> should also test for <code>EAGAIN</code> .
<code>EBADF</code>	<code>fd</code> is not a valid open descriptor.  <code>cmd</code> is <code>F_SETLK</code> or <code>F_SETLKW</code> and the process does not have the appropriate read or write permissions on the file.
<code>EDEADLK</code>	<code>cmd</code> is <code>F_SETLKW</code> , the lock is blocked by one from another process, and putting the calling-process to sleep would cause a deadlock.
<code>EFAULT</code>	<code>cmd</code> is <code>F_GETLK</code> , <code>F_SETLK</code> , or <code>F_SETLKW</code> and <code>arg</code> points to an invalid address.
<code>EINTR</code>	<code>cmd</code> is <code>F_SETLKW</code> and a signal interrupted the process while it was waiting for the lock to be granted.
<code>EINVAL</code>	<code>cmd</code> is <code>F_DUPFD</code> and <code>arg</code> is negative or greater than the maximum allowable number (see <code>getdtablesize(2)</code> ).  <code>cmd</code> is <code>F_GETLK</code> , <code>F_SETLK</code> , or <code>F_SETLKW</code> and <code>arg</code> points to invalid data.

**EMFILE** *cmd* is **F\_DUPFD** and the maximum number of open descriptors has been reached.  
**ENOLCK** *cmd* is **F\_SETLK** or **F\_SETLKW** and there are no more file lock entries available.

**SEE ALSO**

**close(2V)**, **execve(2V)**, **flock(2)**, **fork(2V)**, **getdtablesize(2)**, **ioctl(2)**, **open(2V)**, **sigvec(2)**, **lockf(3)**, **fcntl(5)**, **lockd(8C)**

**WARNINGS**

Mandatory record locks are dangerous. If a runaway or otherwise out-of-control process should hold a mandatory lock on a file critical to the system and fail to release that lock, the entire system could hang or crash. For this reason, mandatory record locks may be removed in a future SunOS release. Use advisory record locking whenever possible.

**NOTES**

Advisory locks allow cooperating processes to perform consistent operations on files, but do not guarantee exclusive access. Files can be accessed without advisory files, but inconsistencies may result.

**read(2V)** and **write(2V)** system calls on files are affected by mandatory file and record locks (see **chmod(2V)**).

**BUGS**

File locks obtained by **fcntl()** do not interact with **flock()** locks. They do, however, work correctly with the exclusive locks claimed by **lockf(3)**.

**F\_GETLK** returns **F\_UNLCK** if the requesting process holds the specified lock. Thus, there is no way for a process to determine if it is still holding a specific lock after catching a **SIGLOST** signal.

In a network environment, the value of **l\_pid** returned by **F\_GETLK** is next to useless.

## NAME

**flock** – apply or remove an advisory lock on an open file

## SYNOPSIS

```
#include <sys/file.h>

#define LOCK_SH      1      /* shared lock */
#define LOCK_EX      2      /* exclusive lock */
#define LOCK_NB      4      /* don't block when locking */
#define LOCK_UN      8      /* unlock */

int flock(fd, operation)
int fd, operation;
```

## DESCRIPTION

**flock()** applies or removes an *advisory* lock on the file associated with the file descriptor *fd*. A lock is applied by specifying an *operation* parameter that is the inclusive OR of **LOCK\_SH** or **LOCK\_EX** and, possibly, **LOCK\_NB**. To unlock an existing lock, the *operation* should be **LOCK\_UN**.

Advisory locks allow cooperating processes to perform consistent operations on files, but do not guarantee exclusive access (that is, processes may still access files without using advisory locks, possibly resulting in inconsistencies).

The locking mechanism allows two types of locks: *shared* locks and *exclusive* locks. More than one process may hold a shared lock for a file at any given time, but multiple exclusive, or both shared and exclusive, locks may not exist simultaneously on a file.

A shared lock may be *upgraded* to an exclusive lock, and vice versa, simply by specifying the appropriate lock type; the previous lock will be released and the new lock applied (possibly after other processes have gained and released the lock).

Requesting a lock on an object that is already locked normally causes the caller to block until the lock may be acquired. If **LOCK\_NB** is included in *operation*, then this will not happen; instead the call will fail and the error **EWOULDBLOCK** will be returned.

## NOTES

Locks are on files, not file descriptors. That is, file descriptors duplicated through **dup(2V)** or **fork(2V)** do not result in multiple instances of a lock, but rather multiple references to a single lock. If a process holding a lock on a file forks and the child explicitly unlocks the file, the parent will lose its lock.

Processes blocked awaiting a lock may be awakened by signals.

## RETURN VALUES

**flock()** returns:

- 0        on success.
- 1       on failure and sets **errno** to indicate the error.

## ERRORS

- EBADF**            The argument **fd** is an invalid descriptor.
- EOPNOTSUPP**      The argument **fd** refers to an object other than a file.
- EWOULDBLOCK**     The file is locked and the **LOCK\_NB** option was specified.

## SEE ALSO

**close(2V)**, **dup(2V)**, **execve(2V)**, **fcntl(2V)**, **fork(2V)**, **open(2V)**, **lockf(3)**, **lockd(8C)**

## BUGS

Locks obtained through the **flock()** mechanism are known only within the system on which they were placed. Thus, multiple clients may successfully acquire exclusive locks on the same remote file. If this behavior is not explicitly desired, the **fcntl(2V)** or **lockf(3)** system calls should be used instead; these make use of the services of the **network lock manager** (see **lockd(8C)**).



**NAME**

**fork** – create a new process

**SYNOPSIS**

**int fork()**

**SYSTEM V SYNOPSIS**

**pid\_t fork()**

**DESCRIPTION**

**fork()** creates a new process. The new process (child process) is an exact copy of the calling process except for the following:

- The child process has a unique process ID. The child process ID also does not match any active process group ID.
- The child process has a different parent process ID (the process ID of the parent process).
- The child process has its own copy of the parent's descriptors. These descriptors reference the same underlying objects, so that, for instance, file pointers in file objects are shared between the child and the parent, so that an **lseek(2V)** on a descriptor in the child process can affect a subsequent **read(2V)** or **write(2V)** by the parent. This descriptor copying is also used by the shell to establish standard input and output for newly created processes as well as to set up pipes.
- The child process has its own copy of the parent's open directory streams (see **directory(3V)**). Each open directory stream in the child process shares directory stream positioning with the corresponding directory stream of the parent.
- All *semadj* values are cleared; see **semop(2)**.
- The child processes resource utilizations are set to 0; see **getrlimit(2)**. The *it\_value* and *it\_interval* values for the **ITIMER\_REAL** timer are reset to 0; see **getitimer(2)**.
- The child process's values of **tms\_utime()**, **tms\_stime()**, **tms\_cutime()**, and **tms\_cstime()** (see **times(3V)**) are set to zero.
- File locks (see **fcntl(2V)**) previously set by the parent are not inherited by the child.
- Pending alarms (see **alarm(3V)**) are cleared for the child process.
- The set of signals pending for the child process is cleared (see **sigvec(2)**).

**RETURN VALUES**

On success, **fork()** returns 0 to the child process and returns the process ID of the child process to the parent process. On failure, **fork()** returns -1 to the parent process, sets **errno** to indicate the error, and no child process is created.

**ERRORS**

**fork()** will fail and no child process will be created if one or more of the following are true:

<b>EAGAIN</b>	The system-imposed limit on the total number of processes under execution would be exceeded. This limit is determined when the system is generated.  The system-imposed limit on the total number of processes under execution by a single user would be exceeded. This limit is determined when the system is generated.
<b>ENOMEM</b>	There is insufficient swap space for the new process.

**SEE ALSO**

**execve(2V)**, **getitimer(2)**, **getrlimit(2)**, **lseek(2V)**, **read(2V)**, **semop(2)**, **wait(2V)**, **write(2V)**

**NAME**

**fsync** – synchronize a file's in-core state with that on disk

**SYNOPSIS**

**int fsync(*fd*)**

**int *fd*;**

**DESCRIPTION**

**fsync()** moves all modified data and attributes of *fd* to a permanent storage device: all in-core modified copies of buffers for the associated file have been written to a disk when the call returns. Note: this is different than **sync(2)** which schedules disk I/O for all files (as though an **fsync()** had been done on all files) but returns before the I/O completes.

**fsync()** should be used by programs which require a file to be in a known state; for example, a program which contains a simple transaction facility might use it to ensure that all modifications to a file or files caused by a transaction were recorded on disk.

**RETURN VALUES**

**fsync()** returns:

0        on success.

-1       on failure and sets **errno** to indicate the error.

**ERRORS**

**EBADF**        *fd* is not a valid descriptor.

**EINVAL**       *fd* refers to a socket, not a file.

**EIO**         An I/O error occurred while reading from or writing to the file system.

**SEE ALSO**

**cron(8), sync(2)**

**NAME**

getaudit, setaudit – get and set user audit identity

**SYNOPSIS**

**int** getaudit()

**int** setaudit(auid)

**int** auid;

**DESCRIPTION**

The **getaudit()** system call returns the audit user ID for the current process. This value is initially set at login time and inherited by all child processes. This value does not change when the real/effective user IDs change, so it can be used to identify the logged-in user, even when running a setuid program. The audit user ID governs audit decisions for a process.

The **setaudit()** system call sets the audit user ID for the current process. Only the super-user may successfully execute these calls.

**RETURN VALUES**

**getaudit()** returns the audit user ID of the current process on success. On failure, it returns -1 and sets **errno** to indicate the error.

**setaudit()** returns:

0        on success.

-1       on failure and sets **errno** to indicate the error.

**ERRORS**

**EINVAL**        The parameter *auid* is not a valid UID.

**EPERM**         The process's effective user ID is not super-user.

**SEE ALSO**

**getuid(2V)**, **setuseraudit(2)**, **audit(8)**

**NAME**

getdents – gets directory entries in a filesystem independent format

**SYNOPSIS**

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

int getdents(fd, buf, nbytes)
int fd;
char *buf;
int nbytes;
```

**DESCRIPTION**

getdents() attempts to put directory entries from the directory referenced by the file descriptor *fd* into the buffer pointed to by *buf*, in a filesystem independent format. Up to *nbytes* bytes of data will be transferred.

The data in the buffer is a series of *dirent* structures each containing the following entries:

```
off_t      d_off;
u_long     d_fileno;
u_short    d_reclen;
u_short    d_namlen;
char       d_name[MAXNAMLEN + 1]; /* see below */
```

The *d\_off* entry contains a value which is interpretable only by the filesystem that generated it. It may be supplied as an offset to *lseek(2V)* to find the entry following the current one in a directory. The *d\_fileno* entry is a number which is unique for each distinct file in the filesystem. Files that are linked by hard links (see *link(2V)*) have the same *d\_fileno*. The *d\_reclen* entry is the length, in bytes, of the directory record. The *d\_name* entry contains a null terminated file name. The *d\_namlen* entry specifies the length of the file name. Thus the actual size of *d\_name* may vary from 1 to *MAXNAMLEN+1*.

The structures are not necessarily tightly packed. The *d\_reclen* entry may be used as an offset from the beginning of a *dirent* structure to the next structure, if any.

Upon return, the actual number of bytes transferred is returned. The current position pointer associated with *fd* is set to point to the directory entry following the last one returned. The pointer is not necessarily incremented by the number of bytes returned by *getdents()*. If the value returned is zero, the end of the directory has been reached. The current position pointer may be set and retrieved by *lseek(2V)*. It is not safe to set the current position pointer to any value other than a value previously returned by *lseek(2V)*, or the value of a *d\_off* entry in a *dirent* structure returned by *getdents()*, or zero.

**RETURN VALUES**

getdents() returns the number of bytes actually transferred on success. On failure, it returns -1 and sets *errno* to indicate the error.

**ERRORS**

EBADF	<i>fd</i> is not a valid file descriptor open for reading.
EFAULT	<i>buf</i> points outside the allocated address space.
EINTR	A read from a slow device was interrupted before any data arrived by the delivery of a signal.
EINVAL	<i>nbytes</i> is not large enough for one directory entry.
ENOTDIR	The file referenced by <i>fd</i> is not a directory.
EIO	An I/O error occurred while reading from or writing to the file system.

**SEE ALSO**

*link(2V)*, *lseek(2V)*, *open(2V)*, *directory(3V)*

**NOTES**

It is strongly recommended, for portability reasons, that programs that deal with directory entries use the **directory(3V)** interface rather than directly calling **getdents()**.

## NAME

`getdirentries` – gets directory entries in a filesystem independent format

## SYNOPSIS

```
int getdirentries(fd, buf, nbytes, basep)
int fd;
char *buf;
int nbytes;
long *basep;
```

## DESCRIPTION

This system call is now obsolete. It is superseded by the `getdents(2)` system call, which returns directory entries in a new format specified in `<sys/dirent.h>`. The file, `<sys/dir.h>`, has also been modified to use the new directory entry format. Programs which currently call `getdirentries()` should be modified to use the new system call and the new include file `dirent.h` or, preferably, to use the `directory(3V)` library routines. The `getdirentries()` system call is retained in the current SunOS release only for purposes of backwards binary compatibility and will be removed in a future major release.

`getdirentries()` attempts to put directory entries from the directory referenced by the file descriptor `fd` into the buffer pointed to by `buf`, in a filesystem independent format. Up to `nbytes` bytes of data will be transferred. `nbytes` must be greater than or equal to the block size associated with the file, see `stat(2V)`. Sizes less than this may cause errors on certain filesystems.

The data in the buffer is a series of structures each containing the following entries:

```
unsigned long d_fileno;
unsigned short d_reclen;
unsigned short d_namlen;
char d_name[MAXNAMELEN + 1]; /* see below */
```

The `d_fileno` entry is a number which is unique for each distinct file in the filesystem. Files that are linked by hard links (see `link(2V)`) have the same `d_fileno`. The `d_reclen` entry is the length, in bytes, of the directory record. The `d_name` entry contains a null terminated file name. The `d_namlen` entry specifies the length of the file name. Thus the actual size of `d_name` may vary from 2 to `MAXNAMELEN+1`.

The structures are not necessarily tightly packed. The `d_reclen` entry may be used as an offset from the beginning of a `direct` structure to the next structure, if any.

Upon return, the actual number of bytes transferred is returned. The current position pointer associated with `fd` is set to point to the next block of entries. The pointer is not necessarily incremented by the number of bytes returned by `getdirentries()`. If the value returned is zero, the end of the directory has been reached. The current position pointer may be set and retrieved by `lseek(2V)`. `getdirentries()` writes the position of the block read into the location pointed to by `basep`. It is not safe to set the current position pointer to any value other than a value previously returned by `lseek(2V)` or a value previously returned in the location pointed to by `basep` or zero.

## RETURN VALUES

`getdirentries()` returns the number of bytes actually transferred on success. On failure, it returns `-1` and sets `errno` to indicate the error.

## ERRORS

EBADF	<code>fd</code> is not a valid file descriptor open for reading.
EFAULT	Either <code>buf</code> or <code>basep</code> points outside the allocated address space.
EINTR	A read from a slow device was interrupted before any data arrived by the delivery of a signal.
EIO	An I/O error occurred while reading from or writing to the file system.

**SEE ALSO**

**getdents(2), link(2V), lseek(2V), open(2V), stat(2V), directory(3V)**

**NAME**

getdomainname, setdomainname – get/set name of current domain

**SYNOPSIS**

```
int getdomainname(name, namelen)
```

```
char *name;
```

```
int namelen;
```

```
int setdomainname(name, namelen)
```

```
char *name;
```

```
int namelen;
```

**DESCRIPTION**

**getdomainname()** returns the name of the domain for the current processor, as previously set by **setdomainname**. The parameter *namelen* specifies the size of the array pointed to by *name*. The returned name is null-terminated unless insufficient space is provided.

**setdomainname()** sets the domain of the host machine to be *name*, which has length *namelen*. This call is restricted to the super-user and is normally used only when the system is bootstrapped.

The purpose of domains is to enable two distinct networks that may have host names in common to merge. Each network would be distinguished by having a different domain name. At the current time, only the Network Information Service (NIS) and **sendmail(8)** make use of domains.

**RETURN VALUES**

**getdomainname()** and **setdomainname()** return:

0        on success.

-1       on failure and set **errno** to indicate the error.

**ERRORS**

**EFAULT**        The *name* parameter gave an invalid address.

In addition to the above, **setdomainname()** will fail if:

**EPERM**        The caller was not the super-user.

**NOTES**

Domain names are limited to 64 characters.

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed.



**NAME**

getdtablesize – get descriptor table size

**SYNOPSIS**

**getdtablesize()**

**DESCRIPTION**

The call **getdtablesize()** returns the current value of the soft limit component of the **RLIMIT\_NOFILE** resource limit. This resource limit governs the maximum value allowable as the index of a newly created descriptor.

**WARNINGS**

**getdtablesize** is implemented as a system call only for binary compatibility with previous releases.

Because of possible intervening **getrlimit(2)** calls affecting **RLIMIT\_NOFILE**, repeated calls to **getdtablesize()** may return different values. Thus it is unwise to cache the return value in an effort to avoid system call overhead, unless it is known that such intervening calls do not occur.

**SEE ALSO**

**close(2V)**, **dup(2V)**, **getrlimit(2)**, **open(2V)**

**NAME**

getgid, getegid – get group identity

**SYNOPSIS**

**int** getgid()

**int** getegid()

**SYSTEM V SYNOPSIS**

**#include** <sys/types.h>

**gid\_t** getgid()

**gid\_t** getegid()

**DESCRIPTION**

**getgid()** returns the real group ID of the current process. **getegid()** returns the effective group ID of the current process.

The GID is specified at login time by the **group** field in the **/etc/passwd** database (see **passwd(5)**).

The effective GID is more transient, and determines additional access permission during execution of a set-GID process, and it is for such processes that **getegid()** is most useful.

**SEE ALSO**

**getuid(2V)**, **setregid(2)**, **setuid(3V)**

## NAME

getgroups, setgroups – get or set supplementary group IDs

## SYNOPSIS

```
int getgroups(gidsetlen, gidset)
```

```
int gidsetlen;
```

```
int gidset[];
```

```
int setgroups(ngroups, gidset)
```

```
int ngroups;
```

```
int gidset[];
```

## SYSTEM V SYNOPSIS

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

```
int getgroups(gidsetlen, gidset)
```

```
int gidsetlen;
```

```
gid_t gidset[];
```

```
int setgroups(ngroups, gidset)
```

```
int ngroups;
```

```
gid_t gidset[];
```

## DESCRIPTION

**getgroups()** gets the current supplementary group IDs of the user process and stores it in the array *gidset*. The parameter *gidsetlen* indicates the number of entries that may be placed in *gidset*. **getgroups()** returns the actual number of entries placed in the *gidset* array. No more than {NGROUPS\_MAX} (see **sysconf(2V)**), will ever be returned. If *gidsetlen* is 0, **getgroups()** returns the number of groups without modifying the *gidset* array.

**setgroups()** sets the supplementary group IDs of the current user process according to the array *gidset*. The parameter *ngroups* indicates the number of entries in the array and must be no more than {NGROUPS\_MAX} (see **sysconf(2V)**).

Only the super-user may set new groups.

## RETURN VALUES

On success, **getgroups()** returns the number of entries placed in the array pointed to by *gidset*. On failure, it returns -1 and sets **errno** to indicate the error.

**setgroups()** returns:

0           on success.

-1           on failure and sets **errno** to indicate the error.

## ERRORS

Either call fails if:

**EFAULT**           The address specified for *gidset* is outside the process address space.

**getgroups()** fails if:

**EINVAL**           The argument *gidsetlen* is smaller than the number of groups in the group set.

**setgroups()** fails if:

**EPERM**            The caller is not the super-user.

## SEE ALSO

**initgroups(3)**

**NAME**

**gethostid** – get unique identifier of current host

**SYNOPSIS**

**gethostid()**

**DESCRIPTION**

**gethostid()** returns the 32-bit identifier for the current host, which should be unique across all hosts. On a Sun workstation, this number is taken from the CPU board's ID PROM.

**SEE ALSO**

**hostid(1)**

**NAME**

gethostname, sethostname – get/set name of current host

**SYNOPSIS**

```
int gethostname(name, namelen)
```

```
char *name;
```

```
int namelen;
```

```
int sethostname(name, namelen)
```

```
char *name;
```

```
int namelen;
```

**DESCRIPTION**

`gethostname()` returns the standard host name for the current processor, as previously set by `sethostname()`. The parameter *namelen* specifies the size of the array pointed to by *name*. The returned name is null-terminated unless insufficient space is provided.

`sethostname()` sets the name of the host machine to be *name*, which has length *namelen*. This call is restricted to the super-user and is normally used only when the system is bootstrapped.

**RETURN VALUES**

`gethostname()` and `sethostname()` return:

0        on success.

-1       on failure and set `errno` to indicate the error.

**ERRORS**

EFAULT        The *name* or *namelen* parameter gave an invalid address.

In addition to the above, `sethostname()` may set `errno` to:

EPERM         The caller was not the super-user.

**SEE ALSO**

`gethostid(2)`

**NOTES**

Host names are limited to `MAXHOSTNAMELEN` (from `<sys/param.h>`) characters, currently 64.

**NAME**

getitimer, setitimer – get/set value of interval timer

**SYNOPSIS**

```
#include <sys/time.h>

int getitimer (which, value)
int which;
struct itimerval *value;

int setitimer (which, value, ovalue)
int which;
struct itimerval *value, *ovalue;
```

**DESCRIPTION**

The system provides each process with three interval timers, defined in `<sys/time.h>`. The `getitimer()` call stores the current value of the timer specified by `which` into the structure pointed to by `value`. The `setitimer()` call sets the value of the timer specified by `which` to the value specified in the structure pointed to by `value`, and if `ovalue` is not a NULL pointer, stores the previous value of the timer in the structure pointed to by `ovalue`.

A timer value is defined by the `itimerval` structure, which includes the following members:

```
struct timeval it_interval; /* timer interval */
struct timeval it_value;   /* current value */
```

If `it_value` is non-zero, it indicates the time to the next timer expiration. If `it_interval` is non-zero, it specifies a value to be used in reloading `it_value` when the timer expires. Setting `it_value` to zero disables a timer; however, `it_value` and `it_interval` must still be initialized. Setting `it_interval` to zero causes a timer to be disabled after its next expiration (assuming `it_value` is non-zero).

Time values smaller than the resolution of the system clock are rounded up to this resolution.

The three timers are:

<b>ITIMER_REAL</b>	Decrements in real time. A SIGALRM signal is delivered when this timer expires.
<b>ITIMER_VIRTUAL</b>	Decrements in process virtual time. It runs only when the process is executing. A SIGVTALRM signal is delivered when it expires.
<b>ITIMER_PROF</b>	Decrements both in process virtual time and when the system is running on behalf of the process. It is designed to be used by interpreters in statistically profiling the execution of interpreted programs. Each time the ITIMER_PROF timer expires, the SIGPROF signal is delivered. Because this signal may interrupt in-progress system calls, programs using this timer must be prepared to restart interrupted system calls.

**RETURN VALUES**

`getitimer()` and `setitimer()` return:

0	on success.
-1	on failure and set <code>errno</code> to indicate the error.

**ERRORS**

The possible errors are:

<b>EFAULT</b>	The <code>value</code> or <code>ovalue</code> parameter specified a bad address.
<b>EINVAL</b>	The <code>value</code> parameter specified a time that was too large to be handled.

**SEE ALSO**

`sigvec(2)`, `gettimeofday(2)`

**NOTES**

Three macros for manipulating time values are defined in `<sys/time.h>`. **timerclear** sets a time value to zero, **timerisset** tests if a time value is non-zero, and **timercmp** compares two time values (beware that `>=` and `<=` do not work with this macro).

## NAME

getmsg – get next message from a stream

## SYNOPSIS

```
#include <stropts.h>

int getmsg(fd, ctlptr, dataptr, flags)
int fd;
struct strbuf *ctlptr;
struct strbuf *dataptr;
int *flags;
```

## DESCRIPTION

`getmsg()` retrieves the contents of a message (see `intro(2)`) located at the **stream head** read queue from a STREAMS file, and places the contents into user specified buffer(s). The message must contain either a data part, a control part or both. The data and control parts of the message are placed into separate buffers, as described below. The semantics of each part is defined by the STREAMS module that generated the message.

*fd* specifies a file descriptor referencing an open stream. *ctlptr* and *dataptr* each point to a **strbuf** structure that contains the following members:

```
int maxlen;    /* maximum buffer length */
int len;       /* length of data */
char *buf;     /* ptr to buffer */
```

where *buf* points to a buffer in which the data or control information is to be placed, and *maxlen* indicates the maximum number of bytes this buffer can hold. On return, *len* contains the number of bytes of data or control information actually received, or is 0 if there is a zero-length control or data part, or is -1 if no data or control information is present in the message. *flags* may be set to the values 0 or `RS_HIPRI` and is used as described below.

*ctlptr* is used to hold the control part from the message and *dataptr* is used to hold the data part from the message. If *ctlptr* (or *dataptr*) is a NULL pointer or the *maxlen* field is -1, the control (or data) part of the message is not processed and is left on the **stream head** read queue and *len* is set to -1. If the *maxlen* field is set to 0 and there is a zero-length control (or data) part, that zero-length part is removed from the read queue and *len* is set to 0. If the *maxlen* field is set to 0 and there are more than zero bytes of control (or data) information, that information is left on the read queue and *len* is set to 0. If the *maxlen* field in *ctlptr* or *dataptr* is less than, respectively, the control or data part of the message, *maxlen* bytes are retrieved. In this case, the remainder of the message is left on the **stream head** read queue and a non-zero return value is provided, as described below under RETURN VALUES. If information is retrieved from a **priority** message, *flags* is set to `RS_HIPRI` on return.

By default, `getmsg()` processes the first priority or non-priority message available on the **stream head** read queue. However, a process may choose to retrieve only priority messages by setting *flags* to `RS_HIPRI`. In this case, `getmsg()` will only process the next message if it is a priority message.

If `O_NDELAY` has not been set, `getmsg()` blocks until a message, of the type(s) specified by *flags* (priority or either), is available on the **stream head** read queue. If `O_NDELAY` has been set and a message of the specified type(s) is not present on the read queue, `getmsg()` fails and sets `errno` to `EAGAIN`.

If a hangup occurs on the stream from which messages are to be retrieved, `getmsg()` will continue to operate normally, as described above, until the **stream head** read queue is empty. Thereafter, it will return 0 in the *len* fields of *ctlptr* and *dataptr*.



**RETURN VALUES**

`getmsg()` returns a non-negative value on success:

0	A full message was read successfully.
MORECTL	More control information is waiting for retrieval. Subsequent <code>getmsg()</code> calls will retrieve the rest of the message.
MOREDATA	More data are waiting for retrieval. Subsequent <code>getmsg()</code> calls will retrieve the rest of the message.
MORECTL   MOREDATA	Both types of information remain.

On failure, `getmsg()` returns `-1` and sets `errno` to indicate the error.

**ERRORS**

EAGAIN	The <code>O_NDELAY</code> flag is set, and no messages are available.
EBADF	<code>fd</code> is not a valid file descriptor open for reading.
EBADMSG	The queued message to be read is not valid for <code>getmsg()</code> .
EFAULT	<code>ctlptr</code> , <code>dataptr</code> , or <code>flags</code> points to a location outside the allocated address space.
EINTR	A signal was caught during the <code>getmsg()</code> system call.
EINVAL	An illegal value was specified in <code>flags</code> . The <code>stream</code> referenced by <code>fd</code> is linked under a multiplexor.
ENOSTR	A <code>stream</code> is not associated with <code>fd</code> .

A `getmsg()` can also fail if a STREAMS error message had been received at the `stream head` before the call to `getmsg()`. The error returned is the value contained in the STREAMS error message.

**SEE ALSO**

`intro(2)`, `poll(2)`, `putmsg(2)`, `read(2V)`, `write(2V)`

**NAME**

getpagesize – get system page size

**SYNOPSIS**

**int** getpagesize()

**DESCRIPTION**

getpagesize() returns the number of bytes in a page. Page granularity is the granularity of many of the memory management calls.

The page size is a *system* page size and may not be the same as the underlying hardware page size.

**SEE ALSO**

pagesize(1), brk(2)

**NAME**

getpeername – get name of connected peer

**SYNOPSIS**

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

**DESCRIPTION**

getpeername() returns the name of the peer connected to socket *s*. The **int** pointed to by the *namelen* parameter should be initialized to indicate the amount of space pointed to by *name*. On return it contains the actual size of the name returned (in bytes). The name is truncated if the buffer provided is too small.

**DIAGNOSTICS**

A 0 is returned if the call succeeds, -1 if it fails.

**ERRORS**

EBADF	The argument <i>s</i> is not a valid descriptor.
EFAULT	The <i>name</i> parameter points to memory not in a valid part of the process address space.
ENOBUFS	Insufficient resources were available in the system to perform the operation.
ENOTCONN	The socket is not connected.
ENOTSOCK	The argument <i>s</i> is a file, not a socket.

**SEE ALSO**

accept(2), bind(2), getsockname(2), socket(2)

**NAME**

**getpgrp, setpgrp** – return or set the process group of a process

**SYNOPSIS**

```
int getpgrp(pid)
int pid;
```

```
int setpgrp(pid, pgrp)
int pgrp;
int pid;
```

**SYSTEM V SYNOPSIS**

```
int getpgrp()
int setpgrp()
```

**DESCRIPTION**

**getpgrp()** returns the process group of the process indicated by *pid*. If *pid* is zero, then the call applies to the calling process.

Process groups are used for distribution of signals, and by terminals to arbitrate requests for their input. Processes that have the same process group as the terminal run in the foreground and may read from the terminal, while others block with a signal when they attempt to read.

This call is thus used by programs such as **cs(1)** to create process groups in implementing job control. The **TIOCGPRP** and **TIOCSPGRP** calls described in **termio(4)** are used to get/set the process group of the control terminal.

**setpgrp()** sets the process group of the specified process, (*pid*) to the process group specified by *pgrp*. If *pid* is zero, then the call applies to the current (calling) process. If *pgrp* is zero and *pid* refers to the calling process, **setpgrp()** behaves identically to **setsid(2V)**.

If the effective user ID of the calling process is not super-user, then the process to be affected must have the same effective user ID as that of the calling process or be a member of the same session as the calling process.

**SYSTEM V DESCRIPTION**

**getpgrp()** returns the process group of the calling process.

**setpgrp()** behaves identically to **setsid()**.

**RETURN VALUES**

**getpgrp()** returns the process group of the indicated process on success. On failure, it returns **-1** and sets **errno** to indicate the error.

**setpgrp()** returns:

**0** on success.

**-1** on failure and sets **errno** to indicate the error.

**SYSTEM V RETURN VALUES**

**getpgrp()** returns the process group of the calling process on success.

**ERRORS**

**setpgrp()** fails, and the process group is not altered when one of the following occurs:

**EACCES** The value of *pid* matches the process ID of a child process of the calling process and the child process has successfully executed one of the **exec()** functions.

**EINVAL** The value of *pgrp* is less than zero or is greater than **MAXPID**, the maximum process ID as defined in **<sys/param.h>**.

**EPERM**

The process indicated by *pid* is a session leader.

The value of *pid* is valid but matches the process ID of a child process of the calling process and the child process is not in the same session as the calling process.

The value of *pgrp* does not match the process ID of the process indicated by *pid* and there is no process with a process group ID that matches the value of *pgrp* in the same session as the calling process.

The requested process has a different effective user ID from that of the calling process and is not a descendent of the calling process.

The calling process is already a process group leader

The process ID of the calling process equals the process group ID of a different process.

**ESRCH**

The value of *pid* does not match the process ID of the calling process or of a child process of the calling process.

The requested process does not exist.

**SEE ALSO**

**csh(1), intro(2), execve(2V), fork(2V), getpid(2V), getuid(2V), kill(2V), setpgid(2V), signal(3V), termio(4)**

**NAME**

**getpid, getppid** – get process identification

**SYNOPSIS**

**int** getpid()

**int** getppid()

**SYSTEM V SYNOPSIS**

**#include** <sys/types.h>

**pid\_t** getpid()

**pid\_t** getppid()

**DESCRIPTION**

**getpid()** returns the process ID of the current process. Most often it is used to generate uniquely-named temporary files.

**getppid()** returns the process ID of the parent of the current process.

**SEE ALSO**

**gethostid(2)**

## NAME

getpriority, setpriority – get/set process nice value

## SYNOPSIS

```
#include <sys/time.h>
#include <sys/resource.h>

int getpriority(which, who)
int which, who;

int setpriority(which, who, niceval)
int which, who, niceval;
```

## DESCRIPTION

The nice value of a process, process group, or user, as indicated by *which* and *who* is obtained with the `getpriority()` call and set with the `setpriority()` call. Process nice values can range from  $-20$  through  $19$ . The default nice value is  $0$ ; lower nice values cause more favorable scheduling.

*which* is one of `PRIO_PROCESS`, `PRIO_PGRP`, or `PRIO_USER`, and *who* is interpreted relative to *which* (a process identifier for `PRIO_PROCESS`, process group identifier for `PRIO_PGRP`, and a user ID for `PRIO_USER`). A zero value of *who* denotes the current process, process group, or user.

The `getpriority()` call returns the lowest numerical nice value of any of the specified processes. The `setpriority()` call sets the nice values of all of the specified processes to the value specified by *niceval*. If *niceval* is less than  $-20$ , a value of  $-20$  is used; if it is greater than  $19$ , a value of  $19$  is used. Only the super-user may use negative nice values.

## RETURN VALUES

Since `getpriority()` can legitimately return the value  $-1$ , it is necessary to clear the external variable `errno` prior to the call, then check it afterward to determine if a  $-1$  is an error or a legitimate value.

`setpriority()` returns:

- $0$  on success.
- $-1$  on failure and sets `errno` to indicate the error.

## ERRORS

`getpriority()` and `setpriority()` may set `errno` to:

- `EINVAL` *which* was not one of `PRIO_PROCESS`, `PRIO_PGRP`, or `PRIO_USER`.
- `ESRCH` No process was located using the *which* and *who* values specified.

In addition to the errors indicated above, `setpriority()` may fail with one of the following errors returned:

- `EACCES` The call to `setpriority()` would have changed a process' nice value to a value lower than its current value, and the effective user ID of the process executing the call was not that of the super-user.
- `EPERM` A process was located, but neither its effective nor real user ID matched the effective user ID of the caller, and neither the effective nor the real user ID of the process executing `setpriority()` was super-user.

## SEE ALSO

`nice(1)`, `ps(1)`, `fork(2V)`, `nice(3v)` `renice(8)`

## BUGS

It is not possible for the process executing `setpriority()` to lower any other process down to its current nice value, without requiring super-user privileges.

These system calls are misnamed. They get and set the nice value, not the kernel scheduling priority. `nice(1)` discusses the relationship between nice value and scheduling priority.

## NAME

getrlimit, setrlimit – control maximum system resource consumption

## SYNOPSIS

```
#include <sys/time.h>
#include <sys/resource.h>

int getrlimit(resource, rlp)
int resource;
struct rlimit *rlp;

int setrlimit(resource, rlp)
int resource;
struct rlimit *rlp;
```

## DESCRIPTION

Limits on the consumption of system resources by the current process and each process it creates may be obtained with the `getrlimit()` call, and set with the `setrlimit()` call.

The *resource* parameter is one of the following:

<b>RLIMIT_CPU</b>	the maximum amount of cpu time (in seconds) to be used by each process.
<b>RLIMIT_FSIZE</b>	the largest size, in bytes, of any single file that may be created.
<b>RLIMIT_DATA</b>	the maximum size, in bytes, of the data segment for a process; this defines how far a program may extend its break with the <code>sbrk()</code> (see <code>brk(2)</code> ) system call.
<b>RLIMIT_STACK</b>	the maximum size, in bytes, of the stack segment for a process; this defines how far a program's stack segment may be extended automatically by the system.
<b>RLIMIT_CORE</b>	the largest size, in bytes, of a core file that may be created.
<b>RLIMIT_RSS</b>	the maximum size, in bytes, to which a process's resident set size may grow. This imposes a limit on the amount of physical memory to be given to a process; if memory is tight, the system will prefer to take memory from processes that are exceeding their declared resident set size.
<b>RLIMIT_NOFILE</b>	one more than the maximum value that the system may assign to a newly created descriptor. This limit constrains the number of descriptors that a process may create.

A resource limit is specified as a soft limit and a hard limit. When a soft limit is exceeded a process may receive a signal (for example, if the cpu time is exceeded), but it will be allowed to continue execution until it reaches the hard limit (or modifies its resource limit). The `rlimit` structure is used to specify the hard and soft limits on a resource,

```
struct rlimit {
    int    rlim_cur;    /* current (soft) limit */
    int    rlim_max;    /* hard limit */
};
```

Only the super-user may raise the maximum limits. Other users may only alter `rlim_cur` within the range from 0 to `rlim_max` or (irreversibly) lower `rlim_max`.

An "infinite" value for a limit is defined as `RLIM_INFINITY` (`0x7fffffff`).

Because this information is stored in the per-process information, this system call must be executed directly by the shell if it is to affect all future processes created by the shell; `limit` is thus a built-in command to `csh(1)`.



The system refuses to extend the data or stack space when the limits would be exceeded in the normal way: a `brk()` or `sbrk()` call will fail if the data space limit is reached, or the process will be sent a `SIGSEGV` when the stack limit is reached which will kill the process unless `SIGSEGV` is handled on a separate signal stack (since the stack cannot be extended, there is no way to send a signal!).

A file I/O operation that would create a file that is too large generates a signal `SIGXFSZ`; this normally terminates the process, but may be caught. When the soft CPU time limit is exceeded, a signal `SIGXCPU` is sent to the offending process.

**RETURN VALUES**

`getrlimit()` and `setrlimit()` return:

- 0        on success.
- 1       on failure and set `errno` to indicate the error.

**ERRORS**

`EFAULT`        The address specified by `rlp` was invalid.

`EINVAL`        An invalid *resource* was specified.

In addition to the above, `setrlimit()` may set `errno` to:

`EINVAL`        The new `rlim_cur` exceeds the new `rlim_max`.

`EPERM`        The limit specified would have raised the maximum limit value, and the caller was not the super-user.

**SEE ALSO**

`csch(1)`, `sh(1)`, `brk(2)`, `getdtablesize(2)`, `quotactl(2)`

**BUGS**

There should be `limit` and `unlimit` commands in `sh(1)` as well as in `csch(1)`.

## NAME

getrusage – get information about resource utilization

## SYNOPSIS

```
#include <sys/time.h>
#include <sys/resource.h>

int getrusage(who, rusage)
int who;
struct rusage *rusage;
```

## DESCRIPTION

**getrusage()** returns information about the resources utilized by the current process, or all its terminated child processes. The interpretation for some values reported, such as **ru\_idrss**, are dependent on the clock tick interval. This interval is an implementation dependent value; for example, on Sun-3 systems the clock tick interval is 1/50 of a second, while on Sun-4 systems the clock tick interval is 1/100 of a second.

The *who* parameter is one of **RUSAGE\_SELF** or **RUSAGE\_CHILDREN**. The buffer to which *rusage* points will be filled in with the following structure:

```
struct rusage {
    struct timeval ru_utime;      /* user time used */
    struct timeval ru_stime;      /* system time used */
    int ru_maxrss;               /* maximum resident set size */
    int ru_ixrss;                /* currently 0 */
    int ru_idrss;                /* integral resident set size */
    int ru_isrss;                /* currently 0 */
    int ru_minflt;              /* page faults not requiring physical I/O */
    int ru_majflt;              /* page faults requiring physical I/O */
    int ru_nswap;                /* swaps */
    int ru_inblock;              /* block input operations */
    int ru_oublock;              /* block output operations */
    int ru_msgsnd;               /* messages sent */
    int ru_msrvcv;               /* messages received */
    int ru_nsignals;             /* signals received */
    int ru_nvcsw;                /* voluntary context switches */
    int ru_nivcsw;               /* involuntary context switches */
};
```

The fields are interpreted as follows:

<b>ru_utime</b>	The total amount of time spent executing in user mode. Time is given in seconds and microseconds.
<b>ru_stime</b>	The total amount of time spent executing in system mode. Time is given in seconds and microseconds.
<b>ru_maxrss</b>	The maximum resident set size. Size is given in pages (the size of a page, in bytes, is given by the <b>getpagesize(2)</b> system call). Also, see <b>WARNINGS</b> .
<b>ru_ixrss</b>	Currently returns 0.
<b>ru_idrss</b>	An “integral” value indicating the amount of memory in use by a process while the process is running. This value is the sum of the resident set sizes of the process running when a clock tick occurs. The value is given in pages times clock ticks. Note: it does not take sharing into account. Also, see <b>WARNINGS</b> .

<b>ru_isrss</b>	Currently returns 0.
<b>ru_minflt</b>	The number of page faults serviced which did not require any physical I/O activity. Also, see WARNINGS.
<b>ru_majflt</b>	The number of page faults serviced which required physical I/O activity. This could include page ahead operations by the kernel. Also, see WARNINGS.
<b>ru_nswap</b>	The number of times a process was swapped out of main memory.
<b>ru_inblock</b>	The number of times the file system had to perform input in servicing a <b>read(2V)</b> request.
<b>ru_oublock</b>	The number of times the file system had to perform output in servicing a <b>write(2V)</b> request.
<b>ru_msgsnd</b>	The number of messages sent over sockets.
<b>ru_msgrcv</b>	The number of messages received from sockets.
<b>ru_nsignals</b>	The number of signals delivered.
<b>ru_nvcsw</b>	The number of times a context switch resulted due to a process voluntarily giving up the processor before its time slice was completed (usually to await availability of a resource).
<b>ru_nivcsw</b>	The number of times a context switch resulted due to a higher priority process becoming runnable or because the current process exceeded its time slice.

**RETURN VALUES**

**getrusage()** returns:

- 0 on success.
- 1 on failure and sets **errno** to indicate the error.

**ERRORS**

- EFAULT** The address specified by the *rusage* argument is not in a valid portion of the process's address space.
- EINVAL** The *who* parameter is not a valid value.

**SEE ALSO**

**gettimeofday(2)**, **read(2V)**, **wait(2V)**, **write(2V)**

**WARNINGS**

The numbers **ru\_inblock** and **ru\_oublock** account only for real I/O, and are approximate measures at best. Data supplied by the caching mechanism is charged only to the first process to read and the last process to write the data.

The way resident set size is calculated is an approximation, and could misrepresent the true resident set size.

Page faults can be generated from a variety of sources and for a variety of reasons. The customary cause for a page fault is a direct reference by the program to a page which is not in memory. Now, however, the kernel can generate page faults on behalf of the user, for example, servicing **read(2V)** and **write(2V)** system calls. Also, a page fault can be caused by an absent hardware translation to a page, even though the page is in physical memory.

In addition to hardware detected page faults, the kernel may cause pseudo page faults in order to perform some housekeeping. For example, the kernel may generate page faults, even if the pages exist in physical memory, in order to lock down pages involved in a raw I/O request.

By definition, *major* page faults require physical I/O, while *minor* page faults do not require physical I/O. For example, reclaiming the page from the free list would avoid I/O and generate a minor page fault. More commonly, minor page faults occur during process startup as references to pages which are already in

memory. For example, if an address space faults on some "hot" executable or shared library, this results in a minor page fault for the address space. Also, any one doing a `read(2V)` or `write(2V)` to something that is in the page cache will get a minor page fault(s) as well.

**BUGS**

There is no way to obtain information about a child process which has not yet terminated.

**NAME**

getsockname – get socket name

**SYNOPSIS**

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

**DESCRIPTION**

**getsockname()** returns the current *name* for the specified socket. The *namelen* parameter should be initialized to indicate the amount of space pointed to by *name*. On return it contains the actual size of the name returned (in bytes).

**DIAGNOSTICS**

A 0 is returned if the call succeeds, -1 if it fails.

**ERRORS**

The call succeeds unless:

EBADF	<i>s</i> is not a valid descriptor.
EFAULT	<i>name</i> points to memory not in a valid part of the process address space.
ENOBUFS	Insufficient resources were available in the system to perform the operation.
ENOTSOCK	<i>s</i> is a file, not a socket.

**SEE ALSO**

**bind(2)**, **getpeername(2)**, **socket(2)**

**BUGS**

Names bound to sockets in the UNIX domain are inaccessible; **getsockname()** returns a zero length name.

## NAME

getsockopt, setsockopt – get and set options on sockets

## SYNOPSIS

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

int getsockopt(s, level, optname, optval, optlen)
int s, level, optname;
char *optval;
int *optlen;

int setsockopt(s, level, optname, optval, optlen)
int s, level, optname;
char *optval;
int optlen;
```

## DESCRIPTION

**getsockopt()** and **setsockopt()** manipulate *options* associated with a socket. Options may exist at multiple protocol levels; they are always present at the uppermost “socket” level.

When manipulating socket options the level at which the option resides and the name of the option must be specified. To manipulate options at the “socket” level, *level* is specified as `SOL_SOCKET`. To manipulate options at any other level the protocol number of the appropriate protocol controlling the option is supplied. For example, to indicate that an option is to be interpreted by the TCP protocol, *level* should be set to the protocol number of TCP; see **getprotoent(3N)**.

The parameters *optval* and *optlen* are used to access option values for **setsockopt()**. For **getsockopt()** they identify a buffer in which the value for the requested option(s) are to be returned. For **getsockopt()**, *optlen* is a value-result parameter, initially containing the size of the buffer pointed to by *optval*, and modified on return to indicate the actual size of the value returned. If no option value is to be supplied or returned, *optval* may be supplied as 0.

*optname* and any specified options are passed uninterpreted to the appropriate protocol module for interpretation. The include file `<sys/socket.h>` contains definitions for “socket” level options, described below. Options at other protocol levels vary in format and name; consult the appropriate entries in section (4P).

Most socket-level options take an *int* parameter for *optval*. For **setsockopt()**, the parameter should be non-zero to enable a boolean option, or zero if the option is to be disabled. `SO_LINGER` uses a **struct linger** parameter, defined in `<sys/socket.h>`, which specifies the desired state of the option and the linger interval (see below).

The following options are recognized at the socket level. Except as noted, each may be examined with **getsockopt()** and set with **setsockopt()**.

<code>SO_DEBUG</code>	toggle recording of debugging information
<code>SO_REUSEADDR</code>	toggle local address reuse
<code>SO_KEEPAIVE</code>	toggle keep connections alive
<code>SO_DONTROUTE</code>	toggle routing bypass for outgoing messages
<code>SO_LINGER</code>	linger on close if data present
<code>SO_BROADCAST</code>	toggle permission to transmit broadcast messages
<code>SO_OOBINLINE</code>	toggle reception of out-of-band data in band
<code>SO_SNDBUF</code>	set buffer size for output
<code>SO_RCVBUF</code>	set buffer size for input
<code>SO_TYPE</code>	get the type of the socket (get only)
<code>SO_ERROR</code>	get and clear error on the socket (get only)

`SO_DEBUG` enables debugging in the underlying protocol modules. `SO_REUSEADDR` indicates that the rules used in validating addresses supplied in a **bind(2)** call should allow reuse of local addresses. `SO_KEEPAIVE` enables the periodic transmission of messages on a connected socket. Should the

connected party fail to respond to these messages, the connection is considered broken. A process attempting to write to the socket receives a SIGPIPE signal and the write operation returns an error. By default, a process exits when it receives SIGPIPE. A read operation on the socket returns an error but does not generate SIGPIPE. If the process is waiting in `select(2)` when the connection is broken, `select()` returns true for any read or write events selected for the socket. `SO_DONTROUTE` indicates that outgoing messages should bypass the standard routing facilities. Instead, messages are directed to the appropriate network interface according to the network portion of the destination address.

`SO_LINGER` controls the action taken when unsent messages are queued on socket and a `close(2V)` is performed. If the socket promises reliable delivery of data and `SO_LINGER` is set, the system will block the process on the `close()` attempt until it is able to transmit the data or until it decides it is unable to deliver the information (a timeout period, termed the linger interval, is specified in the `setsockopt()` call when `SO_LINGER` is requested). If `SO_LINGER` is disabled and a `close()` is issued, the system will process the close in a manner that allows the process to continue as quickly as possible.

The option `SO_BROADCAST` requests permission to send broadcast datagrams on the socket. Broadcast was a privileged operation in earlier versions of the system. With protocols that support out-of-band data, the `SO_OOBINLINE` option requests that out-of-band data be placed in the normal data input queue as received; it will then be accessible with `recv()` or `read()` calls without the `MSG_OOB` flag. `SO_SNDBUF` and `SO_RCVBUF` are options to adjust the normal buffer sizes allocated for output and input buffers, respectively. The buffer size may be increased for high-volume connections, or may be decreased to limit the possible backlog of incoming data. The system places an absolute limit on these values. Finally, `SO_TYPE` and `SO_ERROR` are options used only with `getsockopt()`. `SO_TYPE` returns the type of the socket, such as `SOCK_STREAM`; it is useful for servers that inherit sockets on startup. `SO_ERROR` returns any pending error on the socket and clears the error status. It may be used to check for asynchronous errors on connected datagram sockets or for other asynchronous errors.

#### RETURN VALUES

`getsockopt()` and `setsockopt()` return:

- 0 on success.
- 1 on failure and set `errno` to indicate the error.

#### ERRORS

- `EBADF` *s* is not a valid descriptor.
- `EFAULT` The address pointed to by *optval* is not in a valid part of the process address space.
- `ENOPROTOOPT` The option is unknown at the level indicated.
- `ENOTSOCK` *s* is a file, not a socket.

In addition to the above, `getsockopt()` may set `errno` to:

- `EFAULT` *optlen* is not in a valid part of the process address space.

#### SEE ALSO

`ioctl(2)`, `socket(2)`, `getprotoent(3N)`

#### BUGS

Several of the socket options should be handled at lower levels of the system.

## NAME

gettimeofday, settimeofday – get or set the date and time

## SYNOPSIS

```
#include <sys/time.h>

int gettimeofday(tp, tzp)
struct timeval *tp;
struct timezone *tzp;

int settimeofday(tp, tzp)
struct timeval *tp;
struct timezone *tzp;
```

## DESCRIPTION

The system's notion of the current Greenwich time and the current time zone is obtained with the `gettimeofday()` call, and set with the `settimeofday()` call. The current time is expressed in elapsed seconds and microseconds since 00:00 GMT, January 1, 1970 (zero hour). The resolution of the system clock is hardware dependent; the time may be updated continuously, or in "ticks."

`tp` points to a `timeval` structure, which includes the following members:

```
long tv_sec; /* seconds since Jan. 1, 1970 */
long tv_usec; /* and microseconds */
```

If `tp` is a NULL pointer, the current time information is not returned or set.

`tzp` points to a `timezone()` structure, which includes the following members:

```
int tz_minuteswest; /* of Greenwich */
int tz_dsttime; /* type of dst correction to apply */
```

The `timezone()` structure indicates the local time zone (measured in minutes westward from Greenwich), and a flag that indicates the type of Daylight Saving Time correction to apply. Note: this flag does *not* indicate whether Daylight Saving Time is currently in effect.

Also note that the offset of the local time zone from GMT may change over time, as may the rules for Daylight Saving Time correction. The `localtime()` routine (see `ctime(3V)`) obtains this information from a file rather than from `gettimeofday()`. Programs should use `localtime()` to convert dates and times; the `timezone()` structure is filled in by `gettimeofday()` for backward compatibility with existing programs.

The flag indicating the type of Daylight Saving Time correction should have one of the following values (as defined in `<sys/time.h>`):

```
0    DST_NONE: Daylight Savings Time not observed
1    DST_USA: United States DST
2    DST_AUST: Australian DST
3    DST_WET: Western European DST
4    DST_MET: Middle European DST
5    DST_EET: Eastern European DST
6    DST_CAN: Canadian DST
7    DST_GB: Great Britain and Eire DST
8    DST_RUM: Rumanian DST
9    DST_TUR: Turkish DST
10   DST_AUSTALT: Australian-style DST with shift in 1986
```

If `tzp` is a NULL pointer, the time zone information is not returned or set.

Only the super-user may set the time of day or the time zone.



**RETURN VALUES**

**gettimeofday()** returns:

- 0        on success.
- 1       on failure and sets **errno** to indicate the error.

**ERRORS**

- EFAULT**        An argument address referenced invalid memory.
- EPERM**         A user other than the super-user attempted to set the time or time zone.

**SEE ALSO**

**date(1V)**, **adjtime(2)**, **ctime(3V)**

**BUGS**

Time is never correct enough to believe the microsecond values. There should a mechanism by which, at least, local clusters of systems might synchronize their clocks to millisecond granularity.

**NAME**

`getuid`, `geteuid` – get user identity

**SYNOPSIS**

**int** `getuid()`

**int** `geteuid()`

**SYSTEM V SYNOPSIS**

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

**uid\_t** `getuid()`

**uid\_t** `geteuid()`

**DESCRIPTION**

`getuid()` returns the real user ID of the current process, `geteuid()` the effective user ID.

The real user ID identifies the person who is logged in. The effective user ID gives the process different permissions during execution of “set-user-ID” mode processes, which use `getuid()` to determine the real-user-id of the process that invoked them.

**SEE ALSO**

`getgid(2V)`, `setreuid(2)`

**NAME**

`ioctl` – control device

**SYNOPSIS**

```
int ioctl(fd, request, arg)  
int fd, request;  
caddr_t arg;
```

**DESCRIPTION**

`ioctl()` performs a special function on the object referred to by the open descriptor *fd*. The set of functions that may be performed depends on the object that *fd* refers to. For example, many operating characteristics of character special files (for instance, terminals) may be controlled with `ioctl()` requests. The writeups in section 4 discuss how `ioctl()` applies to various objects.

The *request* codes for particular functions are specified in include files specific to objects or to families of objects; the writeups in section 4 indicate which include files specify which *requests*.

For most `ioctl()` functions, *arg* is a pointer to data to be used by the function or to be filled in by the function. Other functions may ignore *arg* or may treat it directly as a data item; they may, for example, be passed an `int` value.

**RETURN VALUES**

`ioctl()` returns 0 on success for most requests. Some specialized requests may return non-zero values on success; see the description of the request in the man page for the object. On failure, `ioctl()` returns -1 and sets `errno` to indicate the error.

**ERRORS**

<code>EBADF</code>	<i>fd</i> is not a valid descriptor.
<code>EFAULT</code>	<i>request</i> requires a data transfer to or from a buffer pointed to by <i>arg</i> , but some part of the buffer is outside the process's allocated space.
<code>EINVAL</code>	<i>request</i> or <i>arg</i> is not valid.
<code>ENOTTY</code>	The specified request does not apply to the kind of object to which the descriptor <i>fd</i> refers.

`ioctl()` will also fail if the object on which the function is being performed detects an error. In this case, an error code specific to the object and the function will be returned.

**SEE ALSO**

`execve(2V)`, `fcntl(2V)`, `filio(4)`, `mtio(4)`, `sockio(4)`, `streamio(4)`, `termio(4)`

**NAME**

kill – send a signal to a process or a group of processes

**SYNOPSIS**

```
#include <signal.h>
```

```
int kill(pid, sig)
```

```
int pid;
```

```
int sig;
```

**SYSTEM V SYNOPSIS**

```
#include <signal.h>
```

```
int kill(pid, sig)
```

```
pid_t pid;
```

```
int sig;
```

**DESCRIPTION**

kill() sends the signal *sig* to a process or a group of processes. The process or group of processes to which the signal is to be sent is specified by *pid*. *sig* may be one of the signals specified in `sigvec(2)`, or it may be 0, in which case error checking is performed but no signal is actually sent. This can be used to check the validity of *pid* or the existence of process *pid*.

The real or effective user ID of the sending process must match the real or saved set-user ID of the receiving process, unless the effective user ID of the sending process is super-user. A single exception is the signal SIGCONT, which may always be sent to any member of the same session as the current process.

In the following discussion, “system processes” are processes, such as processes 0 and 2, that are not running a regular user program.

If *pid* is greater than zero, the signal is sent to the process whose process ID is equal to *pid*. *pid* may equal 1.

If *pid* is 0, the signal is sent to all processes, except system processes and process 1, whose process group ID is equal to the process group ID of the sender; this is a variant of `killpg(2)`.

If *pid* is -1 and the effective user ID of the sender is not super-user, the signal is sent to all processes, except system processes, process 1, and the process sending the signal, whose real or saved set-user ID matches the real or effective ID of the sender.

If *pid* is -1 and the effective user ID of the sender is super-user, the signal is sent to all processes except system processes, process 1, and the process sending the signal.

If *pid* is negative but not -1, the signal is sent to all processes, except system processes, process 1, and the process sending the signal, whose process group ID is equal to the absolute value of *pid*; this is a variant of `killpg(2)`.

Processes may send signals to themselves.

**SYSTEM V DESCRIPTION**

If a signal is sent to a group of processes (as with, if *pid* is 0 or negative), and if the process sending the signal is a member of that group, the signal is sent to that process as well.

The signal SIGKILL cannot be sent to process 1.

**RETURN VALUES**

kill() returns:

0       on success.

-1       on failure and sets `errno` to indicate the error.

**ERRORS**

**kill()** will fail and no signal will be sent if any of the following occur:

**EINVAL**            *sig* was not a valid signal number.

**EPERM**            The effective user ID of the sending process was not super-user, and neither its real nor effective user ID matched the real or saved set-user ID of the receiving process.

**ESRCH**            No process could be found corresponding to that specified by *pid*.

**SYSTEM V ERRORS**

**kill()** will also fail, and no signal will be sent, if the following occurs:

**EINVAL**            *sig* is SIGKILL and *pid* is 1.

**SEE ALSO**

**getpid(2V), killpg(2), getpgrp(2V), sigvec(2), termio(4)**

**NAME**

**killpg** – send signal to a process group

**SYNOPSIS**

```
int killpg(pgrp, sig)  
int pgrp, sig;
```

**DESCRIPTION**

**killpg()** sends the signal *sig* to the process group *pgrp*. See **sigvec(2)** for a list of signals.

The real or effective user ID of the sending process must match the real or saved set-user ID of the receiving process, unless the effective user ID of the sending process is super-user. A single exception is the signal SIGCONT, which may always be sent to any descendant of the current process.

**RETURN VALUES**

**killpg()** returns:

- 0       on success.
- 1       on failure and sets **errno** to indicate the error.

**ERRORS**

**killpg()** will fail and no signal will be sent if any of the following occur:

- EINVAL**       *sig* was not a valid signal number.
- EPERM**       The effective user ID of the sending process was not super-user, and neither its real nor effective user ID matched the real or saved set-user ID of one or more of the target processes.
- ESRCH**       No processes were found in the specified process group.

**SEE ALSO**

**kill(2V)**, **getpgrp(2V)**, **sigvec(2)**

**NAME**

link – make a hard link to a file

**SYNOPSIS**

```
int link(path1, path2)
char *path1, *path2;
```

**DESCRIPTION**

*path1* points to a pathname naming an existing file. *path2* points to a pathname naming a new directory entry to be created. **link()** atomically creates a new link for the existing file and increments the link count of the file by one. **{LINK\_MAX}** (see **pathconf(2V)**) specifies the maximum allowed number of links to the file.

With hard links, both files must be on the same file system. Both the old and the new link share equal access and rights to the underlying object. The super-user may make multiple links to a directory. Unless the caller is the super-user, the file named by *path1* must not be a directory.

Upon successful completion, **link()** marks for update the **st\_ctime** field of the file. Also, the **st\_ctime** and **st\_mtime** fields of the directory that contains the new entry are marked for update.

**RETURN VALUES**

**link()** returns:

- 0 on success.
- 1 on failure and sets **errno** to indicate the error.

**ERRORS**

**link()** will fail and no link will be created if one or more of the following are true:

- |                     |   |
|---------------------|---|
| <b>EACCES</b>       | Search permission is denied for a component of the path prefix pointed to by <i>path1</i> or <i>path2</i> .<br>The requested link requires writing in a directory for which write permission is denied. |
| <b>EDQUOT</b>       | The directory in which the entry for the new link is being placed cannot be extended because the user's quota of disk blocks on the file system containing the directory has been exhausted.            |
| <b>EEXIST</b>       | The link referred to by <i>path2</i> exists.  |
| <b>EFAULT</b>       | One of the path names specified is outside the process's allocated address space.   |
| <b>EIO</b>          | An I/O error occurred while reading from or writing to the file system to make the directory entry.   |
| <b>ELOOP</b>        | Too many symbolic links were encountered in translating the pathname pointed to by <i>path1</i> or <i>path2</i> .   |
| <b>EMLINK</b>       | The number of links to the file named by <i>path1</i> would exceed <b>{LINK_MAX}</b> (see <b>pathconf(2V)</b> ).  |
| <b>ENAMETOOLONG</b> | The length of the path argument exceeds <b>{PATH_MAX}</b> .<br>A pathname component is longer than <b>{NAME_MAX}</b> while <b>{_POSIX_NO_TRUNC}</b> is in effect (see <b>pathconf(2V)</b> ).            |
| <b>ENOENT</b>       | A component of the path prefix pointed to by <i>path1</i> or <i>path2</i> does not exist.<br>The file referred to by <i>path1</i> does not exist.   |
| <b>ENOSPC</b>       | The directory in which the entry for the new link is being placed cannot be extended because there is no space left on the file system containing the directory.  |
| <b>ENOTDIR</b>      | A component of the path prefix of <i>path1</i> or <i>path2</i> is not a directory.  |

EPERM                   The file named by *path1* is a directory and the effective user ID is not super-user.  
EROFS                   The requested link requires writing in a directory on a read-only file system.  
EXDEV                   The link named by *path2* and the file named by *path1* are on different file systems.

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

ENOENT                 *path1* or *path2* points to an empty string.

**SEE ALSO**

**symlink(2), unlink(2V)**



**NAME**

listen – listen for connections on a socket

**SYNOPSIS**

```
int listen(s, backlog)
int s, backlog;
```

**DESCRIPTION**

To accept connections, a socket is first created with `socket(2)`, a backlog for incoming connections is specified with `listen()` and then the connections are accepted with `accept(2)`. The `listen()` call applies only to sockets of type `SOCK_STREAM` or `SOCK_SEQPACKET`.

The *backlog* parameter defines the maximum length the queue of pending connections may grow to. If a connection request arrives with the queue full the client will receive an error with an indication of `ECONNREFUSED`.

**RETURN VALUES**

`listen()` returns:

- 0        on success.
- 1       on failure and sets `errno` to indicate the error.

**ERRORS**

- `EBADF`                *s* is not a valid descriptor.
- `ENOTSOCK`            *s* is not a socket.
- `EOPNOTSUPP`        The socket is not of a type that supports `listen()`.

**SEE ALSO**

`accept(2)`, `connect(2)`, `socket(2)`

**BUGS**

The *backlog* is currently limited (silently) to 5.

## NAME

`lseek`, `tell` – move read/write pointer

## SYNOPSIS

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

off_t lseek(fd, offset, whence)
int fd;
off_t offset;
int whence;

long tell(fd)
int fd;
```

## DESCRIPTION

`lseek()` sets the seek pointer associated with the open file or device referred to by the descriptor *fd* according to the value supplied for *whence*. *whence* must be one of the following constants defined in `<unistd.h>`:

```
SEEK_SET
SEEK_CUR
SEEK_END
```

If *whence* is `SEEK_SET`, the seek pointer is set to *offset* bytes. If *whence* is `SEEK_CUR`, the seek pointer is set to its current location plus *offset*. If *whence* is `SEEK_END`, the seek pointer is set to the size of the file plus *offset*.

Some devices are incapable of seeking. The value of the seek pointer associated with such a device is undefined.

The obsolete function `tell(fd)` is equivalent to `lseek(fd, 0L, SEEK_CUR)`.

## RETURN VALUES

On success, `lseek()` returns the seek pointer location as measured in bytes from the beginning of the file. On failure, it returns `-1` and sets `errno` to indicate the error.

## ERRORS

`lseek()` will fail and the seek pointer will remain unchanged if:

EBADF	<i>fd</i> is not an open file descriptor.
EINVAL	<i>whence</i> is not a proper value. The seek operation would result in an illegal file offset value for the file (for example, a negative file offset for a file other than a character special file).
ESPIPE	<i>fd</i> is associated with a pipe or a socket.

## SEE ALSO

`dup(2V)`, `open(2V)`

## NOTES

Seeking far beyond the end of a file, then writing, may create a gap or “hole”, which occupies no physical space and reads as zeros.

The constants `L_SET`, `L_INCR`, and `L_XTND` are provided as synonyms for `SEEK_SET`, `SEEK_CUR`, and `SEEK_END`, respectively for backward compatibility but they will disappear in a future release. It is unlikely that the underlying constants 0, 1 and 2 will ever change.

## NAME

mctl – memory management control

## SYNOPSIS

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

int mctl(addr, len, function, arg)
caddr_t addr;
size_t len;
int function;
void *arg;
```

## DESCRIPTION

**mctl()** applies a variety of control functions over pages identified by the mappings established for the address range [*addr*, *addr + len*). The function to be performed is identified by the argument *function*. Legitimate functions are defined in `<sys/mman.h>` as follows.

<b>MC_LOCK</b>	Lock the pages in the range in memory. This function is used to support <b>mlock(3)</b> . See the <b>mlock(3)</b> description for semantics and usage. <i>arg</i> is ignored, but must have the value 0.
<b>MC_LOCKAS</b>	Lock the pages in the address space in memory. This function is used to support <b>mlockall(3)</b> . See the <b>mlockall(3)</b> description for semantics and usage. <i>addr</i> and <i>len</i> are ignored but must be 0. <i>arg</i> is an integer built from the flags: <pre>#define MCL_CURRENT    0x1    /* lock current mappings */ #define MCL_FUTURE    0x2    /* lock future mappings */</pre>
<b>MC_SYNC</b>	Synchronize the pages in the range with their backing storage. Optionally invalidate cache copies. This function is used to support <b>msync(3)</b> . See the <b>msync(3)</b> description for semantics and usage. <i>arg</i> is used to represent the <i>flags</i> argument to <b>msync(3)</b> .
<b>MC_UNLOCK</b>	Unlock the pages in the range. This function is used to support <b>mlock(3)</b> . See the <b>mlock(3)</b> description for semantics and usage. <i>arg</i> is ignored and must have the value 0.
<b>MC_UNLOCKAS</b>	Remove address space memory lock, and locks on all current mappings. This function is used to support <b>mlockall(3)</b> . <i>addr</i> and <i>len</i> must have the value 0. <i>arg</i> is ignored and must have the value 0.

## RETURN VALUES

**mctl()** returns:

- 0 on success.
- 1 on failure and sets **errno** to indicate the error.

## ERRORS

<b>EAGAIN</b>	<i>function</i> was <b>MC_LOCK</b> or <b>MC_LOCKAS</b> and some or all of the memory identified by the operation could not be locked due to insufficient system resources.
<b>EINVAL</b>	<i>addr</i> was not a multiple of the page size as returned by <b>getpagesize(2)</b> . <i>addr</i> and/or <i>len</i> did not have the value 0 when <b>MC_LOCKAS</b> or <b>MC_UNLOCKAS</b> were specified. <i>arg</i> was not valid for the function specified.
<b>ENOMEM</b>	Addresses in the range [ <i>addr</i> , <i>addr + len</i> ) are invalid for the address space of a process, or specify one or more pages which are not mapped.
<b>EPERM</b>	The process's effective user ID was not super-user and one of <b>MC_LOCK</b> , <b>MC_LOCKAS</b> , <b>MC_UNLOCK</b> , or <b>MC_UNLOCKAS</b> was specified.

**SEE ALSO**

**madvise(3), mlock(3), mlockall(3), mmap(2), msync(3)**

**NAME**

**mincore** – determine residency of memory pages

**SYNOPSIS**

```
int mincore(addr, len, vec)  
caddr_t addr; int len; result char *vec;
```

**DESCRIPTION**

**mincore()** returns the primary memory residency status of pages in the address space covered by mappings in the range [*addr*, *addr + len*). The status is returned as a char-per-page in the character array referenced by *\*vec* (which the system assumes to be large enough to encompass all the pages in the address range). The least significant bit of each character is set to 1 to indicate that the referenced page is in primary memory, 0 if it is not. The settings of other bits in each character is undefined and may contain other information in the future.

**RETURN VALUES**

**mincore()** returns:

- 0        on success.
- 1       on failure and sets **errno** to indicate the error.

**ERRORS**

**mincore()** will fail if:

- EFAULT**        A part of the buffer pointer to by *vec* is out-of-range or otherwise inaccessible.
- EINVAL**        *addr* is not a multiple of the page size as returned by **getpagesize(2)**.
- ENOMEM**        Addresses in the range [*addr*, *addr + len*) are invalid for the address space of a process, or specify one or more pages which are not mapped.

**SEE ALSO**

**mmap(2)**

## NAME

**mkdir** – make a directory file

## SYNOPSIS

```
int mkdir(path, mode)
char *path;
int mode;
```

## SYSTEM V SYNOPSIS

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

int mkdir(path, mode)
char *path;
mode_t mode;
```

## DESCRIPTION

**mkdir()** creates a new directory file with name *path*. The mode mask of the new directory is initialized from *mode*.

The low-order 9 bits of *mode* (the file access permissions) are modified such that all bits set in the process's file mode creation mask are cleared (see **umask(2V)**).

The set-GID bit of *mode* is ignored. The set-GID bit of the new file is inherited from that of the parent directory.

The directory's owner ID is set to the process's effective user ID.

The directory's group ID is set to either:

- the effective group ID of the process, if the filesystem was not mounted with the BSD file-creation semantics flag (see **mount(2V)**) and the set-GID bit of the parent directory is clear, or
- the group ID of the directory in which the file is created.

Upon successful completion, **mkdir()** marks for update the **st\_atime**, **st\_ctime**, and **st\_mtime** fields of the directory (see **stat(2V)**). The **st\_ctime** and **st\_mtime** fields of the directory's parent directory are also marked for update.

## RETURN VALUES

**mkdir()** returns:

- 0        on success.
- 1       on failure and sets **errno** to indicate the error.

## ERRORS

**mkdir()** will fail and no directory will be created if:

- |        |   |
|--------|---|
| EACCES | Search permission is denied for a component of the path prefix of <i>path</i> .<br>Write permission is denied on the parent directory of the directory to be created.   |
| EDQUOT | The directory in which the entry for the new file is being placed cannot be extended because the user's quota of disk blocks on the file system containing the directory has been exhausted.<br><br>The new directory cannot be created because the user's quota of disk blocks on the file system which will contain the directory has been exhausted.<br><br>The user's quota of inodes on the file system on which the file is being created has been exhausted. |
| EEXIST | The file referred to by <i>path</i> exists.   |
| EFAULT | <i>path</i> points outside the process's allocated address space.   |

EIO	An I/O error occurred while reading from or writing to the file system.
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
EMLINK	The link count of the parent directory would exceed {LINK_MAX} (see <b>pathconf(2V)</b> ).
ENAMETOOLONG	The length of the path argument exceeds {PATH_MAX}. A pathname component is longer than {NAME_MAX} while {_POSIX_NO_TRUNC} is in effect (see <b>pathconf(2V)</b> ).
ENOENT	A component of the path prefix of <i>path</i> does not exist.
ENOSPC	The directory in which the entry for the new file is being placed cannot be extended because there is no space left on the file system containing the directory. The new directory cannot be created because there is no space left on the file system which will contain the directory. There are no free inodes on the file system on which the file is being created.
ENOTDIR	A component of the path prefix of <i>path</i> is not a directory.
EROFS	<i>path</i> The parent directory of the directory to be created resides on a read-only file system.

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

ENOENT            *path* points to a null pathname.

**SEE ALSO**

**chmod(2V)**, **mount(2V)**, **rmdir(2V)**, **stat(2V)**, **umask(2V)**

## NAME

**mknod, mkfifo** – make a special file

## SYNOPSIS

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

int mknod(path, mode, dev)
char *path;
int mode, dev;

int mkfifo(path, mode)
char *path;
mode_t mode;
```

## DESCRIPTION

**mknod()** creates a new file named by the path name pointed to by *path*. The mode of the new file (including file type bits) is initialized from *mode*. The values of the file type bits which are permitted are:

```
#define S_IFCHR      0020000    /* character special */
#define S_IFBLK     0060000    /* block special */
#define S_IFREG     0100000    /* regular */
#define S_IFIFO     0010000    /* FIFO special */
```

Values of *mode* other than those above are undefined and should not be used.

The access permissions of the mode are modified by the process's mode mask (see **umask(2V)**).

The owner ID of the file is set to the effective user ID of the process. The group ID of the file is set to either:

- the effective group ID of the process, if the filesystem was not mounted with the BSD file-creation semantics flag (see **mount(2V)**) and the set-gid bit of the parent directory is clear, or
- the group ID of the directory in which the file is created.

If *mode* indicates a block or character special file, *dev* is a configuration dependent specification of a character or block I/O device. If *mode* does not indicate a block special or character special device, *dev* is ignored.

**mknod()** may be invoked only by the super-user for file types other than FIFO special.

**mkfifo()** creates a new FIFO special file named by the pathname pointed to by *path*. The access permissions of the new FIFO are initialized from *mode*. The access permissions of *mode* are modified by the process's file creation mask, see **umask(2V)**. Bits in *mode* other than the access permissions are ignored.

The FIFO's owner ID is set to the process's effective user ID. The FIFO's group ID is set to the group ID of the directory in which the FIFO is being created or to the process's effective group ID.

Upon successful completion, the **mkfifo()** function marks for update the **st\_atime**, **st\_ctime**, and **st\_mtime** fields of the file. Also, the **st\_ctime** and **st\_mtime** fields of the directory that contains the new entry are marked for update.

## RETURN VALUES

**mknod()** returns:

- 0 on success.
- 1 on failure and sets **errno** to indicate the error.

**mkfifo()** returns:

- 0 on success.
- 1 on failure and sets **errno** to indicate the error. No FIFO is created.



**ERRORS**

**mknod()** fails and the file mode remains unchanged if:

EACCES	Search permission is denied for a component of the path prefix of <i>path</i> .
EDQUOT	The directory in which the entry for the new file is being placed cannot be extended because the user's quota of disk blocks on the file system containing the directory has been exhausted.
EDQUOT	The user's quota of inodes on the file system on which the node is being created has been exhausted.
EEXIST	The file referred to by <i>path</i> exists.
EFAULT	<i>path</i> points outside the process's allocated address space.
EIO	An I/O error occurred while reading from or writing to the file system.
EISDIR	The specified <i>mode</i> would have created a directory.
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
ENAMETOOLONG	The length of the path argument exceeds {PATH_MAX}. A pathname component is longer than {NAME_MAX} (see <code>sysconf(2V)</code> ) while {_POSIX_NO_TRUNC} is in effect (see <code>pathconf(2V)</code> ).
ENOENT	A component of the path prefix of <i>path</i> does not exist.
ENOSPC	The directory in which the entry for the new file is being placed cannot be extended because there is no space left on the file system containing the directory.
ENOSPC	There are no free inodes on the file system on which the file is being created.
ENOTDIR	A component of the path prefix of <i>path</i> is not a directory.
EPERM	An attempt was made to create a file of type other than FIFO special and the process's effective user ID is not super-user.
EROFS	The file referred to by <i>path</i> resides on a read-only file system.

**mkfifo()** may set `errno` to:

EACCES	A component of the path prefix denies search permission.
EEXIST	The named file already exists.
ENAMETOOLONG	The length of the path string exceeds {PATH_MAX}. A pathname component is longer than {NAME_MAX} while {_POSIX_NO_TRUNC} is in effect (see <code>pathconf(2V)</code> ).
ENOENT	A component of the path prefix does not exist. <i>path</i> points to an empty string.
ENOSPC	The directory that would contain the new file cannot be extended. The file system is out of file allocation resources.
ENOTDIR	A component of the path prefix is not a directory.
EROFS	The named file resides on a read-only file system.

**SEE ALSO**

`chmod(2V)`, `execve(2V)`, `pipe(2V)`, `stat(2V)`, `umask(2V)`, `write(2V)`

## NAME

**mmap** – map pages of memory

## SYNOPSIS

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

caddr_t mmap(addr, len, prot, flags, fd, off)
caddr_t addr;
size_t len;
int prot, flags, fd;
off_t off;
```

## DESCRIPTION

**mmap()** establishes a mapping between the process's address space at an address *pa* for *len* bytes to the memory object represented by *fd* at *off* for *len* bytes. The value of *pa* is an implementation-dependent function of the parameter *addr* and values of *flags*, further described below. A successful **mmap()** call returns *pa* as its result. The address ranges covered by [*pa*, *pa* + *len*) and [*off*, *off* + *len*) must be legitimate for the *possible* (not necessarily current) address space of a process and the object in question, respectively.

The mapping established by **mmap()** replaces any previous mappings for the process's pages in the range [*pa*, *pa* + *len*).

**close(2V)** does not unmap pages of the object referred to by a descriptor. Use **munmap(2)** to remove a mapping.

The parameter *prot* determines whether read, write, execute, or some combination of accesses are permitted to the pages being mapped. The protection options are defined in `<sys/mman.h>` as:

```
#define PROT_READ      0x1      /* page can be read */
#define PROT_WRITE     0x2      /* page can be written */
#define PROT_EXEC      0x4      /* page can be executed */
#define PROT_NONE      0x0      /* page can not be accessed */
```

Not all implementations literally provide all possible combinations. **PROT\_WRITE** is often implemented as **PROT\_READ|PROT\_WRITE** and **PROT\_EXEC** as **PROT\_READ|PROT\_EXEC**. However, no implementation will permit a write to succeed where **PROT\_WRITE** has not been set. The behavior of **PROT\_WRITE** can be influenced by setting **MAP\_PRIVATE** in the *flags* parameter, described below.

The parameter *flags* provides other information about the handling of the mapped pages. The options are defined in `<sys/mman.h>` as:

```
#define MAP_SHARED     1        /* Share changes */
#define MAP_PRIVATE    2        /* Changes are private */
#define MAP_TYPE       0xf      /* Mask for type of mapping */
#define MAP_FIXED      0x10     /* Interpret addr exactly */
```

**MAP\_SHARED** and **MAP\_PRIVATE** describe the disposition of write references to the memory object. If **MAP\_SHARED** is specified, write references will change the memory object. If **MAP\_PRIVATE** is specified, the initial write reference will create a private copy of the memory object page and redirect the mapping to the copy. The mapping type is retained across a **fork(2V)**.

**MAP\_FIXED** informs the system that the value of *pa* must be *addr*, exactly. The use of **MAP\_FIXED** is discouraged, as it may prevent an implementation from making the most effective use of system resources.

When `MAP_FIXED` is not set, the system uses *addr* as a hint in an implementation-defined manner to arrive at *pa*. The *pa* so chosen will be an area of the address space which the system deems suitable for a mapping of *len* bytes to the specified object. All implementations interpret an *addr* value of zero as granting the system complete freedom in selecting *pa*, subject to constraints described below. A non-zero value of *addr* is taken to be a suggestion of a process address near which the mapping should be placed. When the system selects a value for *pa*, it will never place a mapping at address 0, nor will it replace any extant mapping, nor map into areas considered part of the potential data or stack "segments".

The parameter *off* is constrained to be aligned and sized according to the value returned by `getpagesize(2)`. When `MAP_FIXED` is specified, the parameter *addr* must also meet these constraints. The system performs mapping operations over whole pages. Thus, while the parameter *len* need not meet a size or alignment constraint, the system will include in any mapping operation any partial page specified by the range [*pa*, *pa* + *len*).

`mmap()` allows [*pa*, *pa* + *len*) to extend beyond the end of the object, both at the time of the `mmap()` and while the mapping persists, for example if the file was created just prior to the `mmap()` and has no contents, or if the file is truncated. Any reference to addresses beyond the end of the object, however, will result in the delivery of a SIGBUS signal.

The system will always zero-fill any partial page at the end of an object. Further, the system will never write out any modified portions of the last page of an object which are beyond its end. References to whole pages following the end of an object will result in a SIGBUS signal. SIGBUS may also be delivered on various filesystem conditions, including quota exceeded errors.

If the process calls `mlockall(3)` with the `MCL_FUTURE` flag, the pages mapped by all future calls to `mmap()` will be locked in memory. In this case, if not enough memory could be locked, `mmap()` fails and sets `errno` to `EAGAIN`.

#### RETURN VALUES

`mmap()` returns the address at which the mapping was placed (*pa*) on success. On failure, it returns `-1` and sets `errno` to indicate the error.

#### ERRORS

<code>EACCES</code>	<i>fd</i> was not open for read and <code>PROT_READ</code> or <code>PROT_EXEC</code> were specified.
	<i>fd</i> was not open for write and <code>PROT_WRITE</code> was specified for a <code>MAP_SHARED</code> type mapping.
<code>EAGAIN</code>	Some or all of the mapping could not be locked in memory.
<code>EBADF</code>	<i>fd</i> was not open.
<code>EINVAL</code>	The arguments <i>addr</i> (if <code>MAP_FIXED</code> was specified) and <i>off</i> were not multiples of the page size as returned by <code>getpagesize(2)</code> .
	The <code>MAP_TYPE</code> field in <i>flags</i> was invalid (neither <code>MAP_PRIVATE</code> nor <code>MAP_SHARED</code> ).
<code>ENODEV</code>	<i>fd</i> referred to an object for which <code>mmap()</code> is meaningless, such as a terminal.
<code>ENOMEM</code>	<code>MAP_FIXED</code> was specified, and the range [ <i>addr</i> , <i>addr</i> + <i>len</i> ) exceeded that allowed for the address space of a process.
	<code>MAP_FIXED</code> was not specified and there was insufficient room in the address space to effect the mapping.
<code>ENXIO</code>	Addresses in the range [ <i>off</i> , <i>off</i> + <i>len</i> ) are invalid for <i>fd</i> .

#### SEE ALSO

`fork(2V)`, `getpagesize(2)`, `mprotect(2)`, `munmap(2)`, `mlockall(3)`

## NAME

mount – mount file system

## SYNOPSIS

```
#include <sys/mount.h>

int mount(type, dir, M_NEWTYPE| flags, data)
char *type;
char *dir;
int flags;
caddr_t data;
```

## SYSTEM V SYNOPSIS

```
int mount(spec, dir, ronly)
char *spec;
char *dir;
int ronly;
```

## DESCRIPTION

**mount()** attaches a file system to a directory. After a successful return, references to directory *dir* will refer to the root directory on the newly mounted file system. *dir* is a pointer to a null-terminated string containing a path name. *dir* must exist already, and must be a directory. Its old contents are inaccessible while the file system is mounted.

**mount()** may be invoked only by the super-user.

The *flags* argument is constructed by the logical OR of the following bits (defined in `<sys/mount.h>`):

<b>M_RDONLY</b>	mount filesystem read-only.
<b>M_NOSUID</b>	ignore set-uid bit on execution.
<b>M_NEWTYPE</b>	this flag must always be set.
<b>M_GRPID</b>	use BSD file-creation semantics (see <code>open(2V)</code> ).
<b>M_REMOUNT</b>	change options on an existing mount.
<b>M_NOSUB</b>	disallow mounts beneath this filesystem.

Physically write-protected and magnetic tape file systems must be mounted read-only or errors will occur when access times are updated, whether or not any explicit write is attempted.

The *type* string indicates the type of the filesystem. *data* is a pointer to a structure which contains the type specific arguments to mount. Below is a list of the filesystem types supported and the type specific arguments to each:

## 4.2

```
struct ufs_args {
    char *fspec; /* Block special file to mount */
};

"lo"
struct lo_args {
    char *fsdir; /* Pathname of directory to mount */
};

"nfs"
#include <nfs/nfs.h>
#include <netinet/in.h>
struct nfs_args {
    struct sockaddr_in *addr; /* file server address */
    fhandle_t *fh; /* File handle to be mounted */
    int flags; /* flags */
};
```

```

    int    wsize;    /* write size in bytes */
    int    rsize;    /* read size in bytes */
    int    timeo;    /* initial timeout in .1 secs */
    int    retrans;  /* times to retry send */
    char   *hostname; /* server's hostname */
    int    acregmin; /* attr cache file min secs */
    int    acregmax; /* attr cache file max secs */
    int    acdirmin; /* attr cache dir min secs */
    int    acdirmax; /* attr cache dir max secs */
    char   *netname; /* server's netname */

};

rfs
struct rfs_args {
    char   *rmtfs    /* name of remote resource */
    struct token {
        int    t_id; /* token id */
        char   t_uname[64]; /* domain.machine name */
    }
    *token; /* Identifier of remote machine */
};

```

**SYSTEM V DESCRIPTION**

**mount()** requests that a file system contained on the block special file identified by *spec* be mounted on the directory identified by *dir*. *spec* and *dir* point to path names. When **mount()** succeeds, subsequent references to the file named by *dir* refer to the root directory on the mounted file system.

The **M\_RDONLY** bit of *rdonly* is used to control write permission on the mounted file system. If the bit is set, writing is not allowed. Otherwise, writing is permitted according to the access permissions of individual files.

**RETURN VALUES**

**mount()** returns:

- 0        on success.
- 1      on failure and sets **errno** to indicate the error.

**ERRORS**

- EACCES**        Search permission is denied for a component of the path prefix of *dir*.
  - EBUSY**        Another process currently holds a reference to *dir*.
  - EFAULT**       *dir* points outside the process's allocated address space.
  - ELOOP**        Too many symbolic links were encountered in translating the path name of *dir*.
  - ENAMETOOLONG**    The length of the path argument exceeds **{PATH\_MAX}**.  
A pathname component is longer than **{NAME\_MAX}** (see **sysconf(2V)**) while **{\_POSIX\_NO\_TRUNC}** is in effect (see **pathconf(2V)**).
  - ENODEV**        The file system type specified by *type* is not valid or is not configured into the system.
  - ENOENT**        A component of *dir* does not exist.
  - ENOTDIR**       The file named by *dir* is not a directory.
  - EPERM**        The caller is not the super-user.
- For a 4.2 file system, **mount()** fails when one of the following occurs:
- EACCES**        Search permission is denied for a component of the path prefix of *fspec*.

EFAULT	<i>fspec</i> points outside the process's allocated address space.
EINVAL	The super block for the file system had a bad magic number or an out of range block size.
EIO	An I/O error occurred while reading from or writing to the file system.
ELOOP	Too many symbolic links were encountered in translating the path name of <i>fspec</i> .
EMFILE	No space remains in the mount table.
ENAMETOOLONG	The length of the path argument exceeds {PATH_MAX}. A pathname component is longer than {NAME_MAX} (see <code>sysconf(2V)</code> ) while {_POSIX_NO_TRUNC} is in effect (see <code>pathconf(2V)</code> ).
ENOENT	A component of <i>fspec</i> does not exist.
ENOMEM	Not enough memory was available to read the cylinder group information for the file system.
ENOTBLK	<i>fspec</i> is not a block device.
ENOTDIR	A component of the path prefix of <i>fspec</i> is not a directory.
ENXIO	The major device number of <i>fspec</i> is out of range (this indicates no device driver exists for the associated hardware).

**SYSTEM V ERRORS**

EBUSY	The device referred to by <i>spec</i> is currently mounted. There are no more mount table entries.
ENOENT	The file referred to by <i>spec</i> or <i>dir</i> does not exist.
ENOTBLK	<i>spec</i> is not a block special device.
ENOTDIR	A component of the path prefix of <i>dir</i> or <i>spec</i> is not a directory.
ENXIO	The device referred to by <i>spec</i> does not exist.

**SEE ALSO**

`unmount(2V)`, `open(2V)`, `lofs(4S)`, `fstab(5)`, `mount(8)`

**BUGS**

Some of the error codes need translation to more obvious messages.

**NAME**

`mprotect` – set protection of memory mapping

**SYNOPSIS**

```
#include <sys/mman.h>

mprotect(addr, len, prot)
caddr_t addr;
int len, prot;
```

**DESCRIPTION**

`mprotect()` changes the access protections on the mappings specified by the range  $[addr, addr + len)$  to be that specified by *prot*. Legitimate values for *prot* are the same as those permitted for `mmap(2)`.

**RETURN VALUES**

`mprotect()` returns:

- 0       on success.
- 1       on failure and sets `errno` to indicate the error.

**ERRORS**

- EACCESS**       *prot* specifies a protection which violates the access permission the process has to the underlying memory object.
- EINVAL**       *addr* is not a multiple of the page size as returned by `getpagesize(2)`.
- ENOMEM**       Addresses in the range  $[addr, addr + len)$  are invalid for the address space of a process, or specify one or more pages which are not mapped.

When `mprotect()` fails for reasons other than `EINVAL`, the protections on some of the pages in the range  $[addr, addr + len)$  will have been changed. If the error occurs on some page at address *addr2*, then the protections of all whole pages in the range  $[addr, addr2)$  have been modified.

**SEE ALSO**

`getpagesize(2)`, `mmap(2)`

## NAME

msgctl – message control operations

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/msg.h>

int msgctl (msqid, cmd, buf)
int msqid, cmd;
struct msqid_ds *buf;
```

## DESCRIPTION

msgctl() provides a variety of message control operations as specified by *cmd*. The following *cmds* are available:

**IPC\_STAT** Place the current value of each member of the data structure associated with *msqid* into the structure pointed to by *buf*. The contents of this structure are defined in *intro(2)*.

**IPC\_SET** Set the value of the following members of the data structure associated with *msqid* to the corresponding value found in the structure pointed to by *buf*:

```
msg_perm.uid
msg_perm.gid
msg_perm.mode /* only low 9 bits */
msg_qbytes
```

This *cmd* can only be executed by a process that has an effective user ID equal to either that of super-user, or to the value of *msg\_perm.cuid* or *msg\_perm.uid* in the data structure associated with *msqid*. Only super-user can raise the value of *msg\_qbytes*.

**IPC\_RMID** Remove the message queue identifier specified by *msqid* from the system and destroy the message queue and data structure associated with it. This *cmd* can only be executed by a process that has an effective user ID equal to either that of super-user, or to the value of *msg\_perm.cuid* or *msg\_perm.uid* in the data structure associated with *msqid*.

In the *msgop(2)* and *msgctl(2)* system call descriptions, the permission required for an operation is given as "[token]", where "token" is the type of permission needed interpreted as follows:

00400	Read by user
00200	Write by user
00060	Read, Write by group
00006	Read, Write by others

Read and Write permissions on a *msqid* are granted to a process if one or more of the following are true:

The effective user ID of the process is super-user.

The effective user ID of the process matches *msg\_perm.[c]uid* in the data structure associated with *msqid* and the appropriate bit of the "user" portion (0600) of *msg\_perm.mode* is set.

The effective user ID of the process does not match *msg\_perm.[c]uid* and the effective group ID of the process matches *msg\_perm.[c]gid* and the appropriate bit of the "group" portion (060) of *msg\_perm.mode* is set.

The effective user ID of the process does not match *msg\_perm.[c]uid* and the effective group ID of the process does not match *msg\_perm.[c]gid* and the appropriate bit of the "other" portion (06) of *msg\_perm.mode* is set.

Otherwise, the corresponding permissions are denied.



**RETURN VALUES**

**msgctl()** returns:

- 0        on success.
- 1       on failure and sets **errno** to indicate the error.

**ERRORS**

- EACCES**        *cmd* is equal to **IPC\_STAT** and **[READ]** operation permission is denied to the calling process (see **intro(2)**).
- EFAULT**        *buf* points to an illegal address.
- EINVAL**        *msqid* is not a valid message queue identifier.  
*cmd* is not a valid command.
- EPERM**        *cmd* is equal to **IPC\_RMID** or **IPC\_SET**. The effective user ID of the calling process is neither super-user, nor the value of **msg\_perm.cuid** or **msg\_perm.uid** in the data structure associated with *msqid*.  
*cmd* is equal to **IPC\_SET**, an attempt is being made to increase to the value of **msg\_qbytes**, and the effective user ID of the calling process is not equal to that of super-user.

**SEE ALSO**

**intro(2)**, **msgget(2)**, **msgop(2)**

## NAME

msgget – get message queue

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/msg.h>

int msgget(key, msgflg)
key_t key;
int msgflg;
```

## DESCRIPTION

msgget() returns the message queue identifier associated with *key*.

A message queue identifier and associated message queue and data structure (see [intro\(2\)](#)) are created for *key()* if one of the following is true:

- *key* is equal to `IPC_PRIVATE`.
- *key* does not already have a message queue identifier associated with it, and (*msgflg* & `IPC_CREAT`) is “true”.

Upon creation, the data structure associated with the new message queue identifier is initialized as follows:

- `msg_perm.cuid`, `msg_perm.uid`, `msg_perm.cgid`, and `msg_perm.gid` are set equal to the effective user ID and effective group ID, respectively, of the calling process.
- The low-order 9 bits of `msg_perm.mode` are set equal to the low-order 9 bits of *msgflg*.
- `msg_qnum`, `msg_lspid`, `msg_lrpid`, `msg_stime`, and `msg_rtime` are set equal to 0.
- `msg_ctime` is set equal to the current time.
- `msg_qbytes` is set equal to the system-wide standard value of the maximum number of bytes allowed on a message queue.

A message queue identifier (*msqid*) is a unique positive integer created by a `msgget(2)` system call. Each *msqid* has a message queue and a data structure associated with it. The data structure is referred to as `msqid_ds()` and contains the following members:

```
struct ipc_perm msg_perm; /* operation permission struct */
ushort msg_qnum;         /* number of msgs on q */
ushort msg_qbytes;      /* max number of bytes on q */
ushort msg_lspid;       /* pid of last msgsnd operation */
ushort msg_lrpid;       /* pid of last msgrcv operation */
time_t msg_stime;       /* last msgsnd time */
time_t msg_rtime;       /* last msgrcv time */
time_t msg_ctime;       /* last change time */
/* Times measured in secs since */
/* 00:00:00 GMT, Jan. 1, 1970 */
```

`msg_perm()` is an `ipc_perm` structure that specifies the message operation permission (see below). This structure includes the following members:

```
ushort cuid;           /* creator user id */
ushort cgid;           /* creator group id */
ushort uid;            /* user id */
ushort gid;            /* group id */
ushort mode;           /* r/w permission */
```

`msg_qnum` is the number of messages currently on the queue. `msg_qbytes` is the maximum number of bytes allowed on the queue. `msg_lspid` is the process ID of the last process that performed a `msgsnd` operation. `msg_lrpid` is the process ID of the last process that performed a `msgrcv` operation. `msg_stime`

is the time of the last *msgsnd* operation, *msg\_rtime* is the time of the last *msgrcv* operation, and *msg\_ctime* is the time of the last *msgctl(2)* operation that changed a member of the above structure.

**RETURN VALUES**

*msgget()* returns A non-negative message queue identifier on success. On failure, it returns -1 and sets *errno* to indicate the error.

**ERRORS**

- |        |  |
|--------|--|
| EACCES | A message queue identifier exists for <i>key</i> , but operation permission (see <i>intro(2)</i> ) as specified by the low-order 9 bits of <i>msgflg</i> would not be granted. |
| EEXIST | A message queue identifier exists for <i>key()</i> but ( ( <i>msgflg</i> & <i>IPC_CREAT</i> ) & ( <i>msgflg</i> & <i>IPC_EXCL</i> )) is "true".                                |
| ENOENT | A message queue identifier does not exist for <i>key()</i> and ( <i>msgflg</i> & <i>IPC_CREAT</i> ) is "false".  |
| ENOSPC | A message queue identifier is to be created but the system-imposed limit on the maximum number of allowed message queue identifiers system wide would be exceeded.             |

**SEE ALSO**

*intro(2)*, *msgctl(2)*, *msgop(2)*

## NAME

msgop, msgsnd, msgrcv – message operations

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/msg.h>

int msgsnd(msqid, msgp, msgsz, msgflg)
int msqid;
struct msgbuf *msgp;
int msgsz, msgflg;

int msgrcv(msqid, msgp, msgsz, msgtyp, msgflg)
int msqid;
struct msgbuf *msgp;
int msgsz;
long msgtyp;
int msgflg;
```

## DESCRIPTION

**msgsnd()** is used to send a message to the queue associated with the message queue identifier specified by *msqid*. [WRITE] (see **msgctl(2)**) *msgp* points to a structure containing the message. This structure is composed of the following members:

```
long    mtype;    /* message type */
char    mtext[1]; /* message text */
```

*mtype* is a positive integer that can be used by the receiving process for message selection (see **msgrcv()** below). *mtext* is any text of length *msgsz* bytes. *msgsz* can range from 0 to a system-imposed maximum.

*msgflg* specifies the action to be taken if one or more of the following are true:

- The number of bytes already on the queue is equal to *msg\_qbytes* (see **intro(2)**).
- The total number of messages on all queues system-wide is equal to the system-imposed limit.

These actions are as follows:

- If (**msgflg** & **IPC\_NOWAIT**) is “true”, the message will not be sent and the calling process will return immediately.
- If (**msgflg** & **IPC\_NOWAIT**) is “false”, the calling process will suspend execution until one of the following occurs:
  - The condition responsible for the suspension no longer exists, in which case the message is sent.
  - *msqid* is removed from the system (see **msgctl(2)**). When this occurs, **errno** is set equal to **EIDRM**, and a value of **-1** is returned.
  - The calling process receives a signal that is to be caught. In this case the message is not sent and the calling process resumes execution in the manner prescribed in **signal(3V)**.

Upon successful completion, the following actions are taken with respect to the data structure associated with *msqid* (see **intro(2)**).

- *msg\_qnum* is incremented by 1.
- *msg\_lspid* is set equal to the process ID of the calling process.
- *msg\_stime* is set equal to the current time.

**msgrcv()** reads a message from the queue associated with the message queue identifier specified by *msqid* and places it in the structure pointed to by *msgp*. [READ] This structure is composed of the following members:

```

long    mtype;    /* message type */
char    mtext[1]; /* message text */

```

*mtype* is the received message's type as specified by the sending process. *mtext* is the text of the message. *msgsz* specifies the size in bytes of *mtext*. The received message is truncated to *msgsz* bytes if it is larger than *msgsz* and (*msgflg* & *MSG\_NOERROR*) is "true". The truncated part of the message is lost and no indication of the truncation is given to the calling process.

*msgtyp* specifies the type of message requested as follows:

- If *msgtyp* is equal to 0, the first message on the queue is received.
- If *msgtyp* is greater than 0, the first message of type *msgtyp* is received.
- If *msgtyp* is less than 0, the first message of the lowest type that is less than or equal to the absolute value of *msgtyp* is received.

*msgflg* specifies the action to be taken if a message of the desired type is not on the queue. These are as follows:

- If (*msgflg* & *IPC\_NOWAIT*) is "true", the calling process will return immediately with a return value of -1 and *errno* set to *ENOMSG*.
- If (*msgflg* & *IPC\_NOWAIT*) is "false", the calling process will suspend execution until one of the following occurs:
  - A message of the desired type is placed on the queue.
  - *msqid* is removed from the system. When this occurs, *errno* is set equal to *EIDRM*, and a value of -1 is returned.
  - The calling process receives a signal that is to be caught. In this case a message is not received and the calling process resumes execution in the manner prescribed in *signal(3V)*.

Upon successful completion, the following actions are taken with respect to the data structure associated with *msqid* (see *intro(2)*).

- *msg\_qnum* is decremented by 1.
- *msg\_lrp* is set equal to the process ID of the calling process.
- *msg\_rtime* is set equal to the current time.

#### RETURN VALUES

*msgsnd()* returns:

- 0 on success.
- 1 on failure and sets *errno* to indicate the error.

*msgrcv()* returns the number of bytes actually placed into *mtext* on success. On failure, it returns -1 and sets *errno* to indicate the error.

#### ERRORS

*msgsnd()* will fail and no message will be sent if one or more of the following are true:

- |        |  |
|--------|--|
| EACCES | Operation permission is denied to the calling process (see <i>intro(2)</i> ).                                      |
| EAGAIN | The message cannot be sent for one of the reasons cited above and ( <i>msgflg</i> & <i>IPC_NOWAIT</i> ) is "true". |
| EFAULT | <i>msgp</i> points to an illegal address.  |
| EIDRM  | The message queue referred to by <i>msqid</i> was removed from the system.   |
| EINTR  | The call was interrupted by the delivery of a signal.  |

- EINVAL**      *msgid* is not a valid message queue identifier.  
                 *mtype* is less than 1.  
                 *msgsz* is less than zero or greater than the system-imposed limit.
- msgrcv()** will fail and no message will be received if one or more of the following are true:
- E2BIG**      *mtext* is greater than *msgsz* and (**msgflg** & **MSG\_NOERROR**) is "false".
- EACCES**      Operation permission is denied to the calling process.
- EFAULT**      *msgp* points to an illegal address.
- EIDRM**      The message queue referred to by *msgid* was removed from the system.
- EINTR**      The call was interrupted by the delivery of a signal.
- EINVAL**      *msgid* is not a valid message queue identifier.  
                 *msgsz* is less than 0.
- ENOMSG**      The queue does not contain a message of the desired type and (**msgtyp** & **IPC\_NOWAIT**) is "true".

**SEE ALSO****intro(2), msgctl(2), msgget(2), signal(3V)**

**NAME**

`msync` – synchronize memory with physical storage

**SYNOPSIS**

```
#include <sys/mman.h>

int msync(addr, len, flags)
caddr_t addr;
int len, flags;
```

**DESCRIPTION**

`msync()` writes all modified copies of pages over the range  $[addr, addr + len)$  to their permanent storage locations. `msync()` optionally invalidates any copies so that further references to the pages will be obtained by the system from their permanent storage locations.

Values for *flags* are defined in `<sys/mman.h>` as:

```
#define MS_ASYNC      0x1    /* Return immediately */
#define MS_INVALIDATE 0x2    /* Invalidate mappings */
```

and are used to control the behavior of `msync()`. One or more flags may be specified in a single call.

`MS_ASYNC` returns `msync()` immediately once all I/O operations are scheduled; normally, `msync()` will not return until all I/O operations are complete. `MS_INVALIDATE` invalidates all cached copies of data from memory objects, requiring them to be re-obtained from the object's permanent storage location upon the next reference.

`msync()` should be used by programs which require a memory object to be in a known state, for example in building transaction facilities.

**RETURN VALUES**

`msync()` returns:

- 0        on success.
- 1       on failure and sets `errno` to indicate the error.

**ERRORS**

- EINVAL**        *addr* is not a multiple of the current page size.  
                   *len* is negative.  
                   One of the flags `MS_ASYNC` or `MS_INVALID` is invalid.
- EIO**            An I/O error occurred while reading from or writing to the file system.
- ENOMEM**        Addresses in the range  $[addr, addr + len)$  are outside the valid range for the address space of a process.

**NAME**

**munmap** – unmap pages of memory.

**SYNOPSIS**

```
#include <sys/mman.h>
```

```
int munmap(addr, len)
```

```
caddr_t addr;
```

```
int len;
```

**DESCRIPTION**

**munmap()** removes the mappings for pages in the range [*addr*, *addr + len*). Further references to these pages will result in the delivery of a SIGSEGV signal to the process, unless these pages are considered part of the “data” or “stack” segments.

**brk()** and **mmap()** often perform implicit **munmap**'s.

**RETURN VALUES**

**munmap()** returns:

0        on success.

-1       on failure and sets **errno** to indicate the error.

**ERRORS**

**EINVAL**        *addr* is not a multiple of the page size as returned by **getpagesize(2)**.

Addresses in the range [*addr*, *addr + len*) are outside the valid range for the address space of a process.

**SEE ALSO**

**brk(2)**, **getpagesize(2)**, **mmap(2)**



**NAME**

**nfssvc, async\_daemon** – NFS daemons

**SYNOPSIS**

```
nfssvc(sock)  
int sock;  
async_daemon()
```

**DESCRIPTION**

**nfssvc()** starts an NFS daemon listening on socket *sock*. The socket must be **AF\_INET**, and **SOCK\_DGRAM** (protocol **UDP/IP**). The system call will return only if the socket is invalid.

**async\_daemon()** implements the NFS daemon that handles asynchronous I/O for an NFS client. This system call never returns.

Both system calls result in kernel-only processes with user memory discarded.

**SEE ALSO**

**mountd(8C)**

**BUGS**

There should be a way to dynamically create kernel-only processes instead of having to make system calls from userland to simulate this.

## NAME

open – open or create a file for reading or writing

## SYNOPSIS

```
#include <fcntl.h>

int open(path, flags[ , mode ] )
char *path;
int flags;
int mode;
```

## SYSTEM V SYNOPSIS

```
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>

int open(path, flags[ , mode ] )
char *path;
int flags;
mode_t mode;
```

## DESCRIPTION

*path* points to the pathname of a file. **open()** opens the named file for reading and/or writing, as specified by the *flags* argument, and returns a descriptor for that file. The *flags* argument may indicate the file is to be created if it does not already exist (by specifying the **O\_CREAT** flag), in which case the file is created with mode *mode* as described in **chmod(2V)** and modified by the process' umask value (see **umask(2V)**). If the path is an empty string, the kernel maps this empty pathname to '.', the current directory. *flags* values are constructed by ORing flags from the following list (one and only one of the first three flags below must be used):

- O\_RDONLY** Open for reading only.
- O\_WRONLY** Open for writing only.
- O\_RDWR** Open for reading and writing.
- O\_NDELAY** When opening a FIFO (named pipe – see **mknod(2V)**) with **O\_RDONLY** or **O\_WRONLY** set:
  - If **O\_NDELAY** is set:
    - An **open()** for reading-only returns without delay. An **open()** for writing-only returns an error if no process currently has the file open for reading.
  - If **O\_NDELAY** is clear:
    - A call to **open()** for reading-only blocks until a process opens the file for writing. A call to **open()** for writing-only blocks until a process opens the file for reading.
- When opening a file associated with a communication line:
  - If **O\_NDELAY** is set:
    - A call to **open()** returns without waiting for carrier.
  - If **O\_NDELAY** is clear:
    - A call to **open()** blocks until carrier is present.
- O\_NOCTTY** When this flag is set, and *path* refers to a terminal device, **open()** prevents the terminal device from becoming the controlling terminal for the process.
- O\_NONBLOCK** Same as **O\_NDELAY** above.

- O\_SYNC** When opening a regular file, this flag affects subsequent writes. If set, each `write(2V)` will wait for both the file data and file status to be physically updated.
- O\_APPEND** If set, the seek pointer will be set to the end of the file prior to each write.
- O\_CREAT** If the file exists, this flag has no effect. Otherwise, the file is created, and the owner ID of the file is set to the effective user ID of the process. The group ID of the file is set to either:
- the effective group ID of the process, if the filesystem was not mounted with the BSD file-creation semantics flag (see `mount(2V)`) and the set-gid bit of the parent directory is clear, or
  - the group ID of the directory in which the file is created.
- The low-order 12 bits of the file mode are set to the value of *mode*, modified as follows (see `creat(2V)`):
- All bits set in the file mode creation mask of the process are cleared. See `umask(2V)`.
  - The “save text image after execution” bit of the mode is cleared. See `chmod(2V)`.
  - The “set group ID on execution” bit of the mode is cleared if the effective user ID of the process is not super-user and the process is not a member of the group of the created file.
- O\_TRUNC** If the file exists and is a regular file, and the file is successfully opened `O_RDWR` or `O_WRONLY`, its length is truncated to zero and the mode and owner are unchanged. `O_TRUNC` has no effect on FIFO special files or directories.
- O\_EXCL** If `O_EXCL` and `O_CREAT` are set, `open()` will fail if the file exists. This can be used to implement a simple exclusive access locking mechanism.

The seek pointer used to mark the current position within the file is set to the beginning of the file.

The new descriptor is set to remain open across `execve(2V)` system calls; see `close(2V)` and `fcntl(2V)`.

There is a system enforced limit on the number of open file descriptors per process, whose value is returned by the `getdtablesize(2)` call.

If `O_CREAT` is set and the file did not previously exist, upon successful completion, `open()` marks for update the `st_atime`, `st_ctime`, and `st_mtime` fields of the file and the `st_ctime` and `st_mtime` fields of the parent directory.

If `O_TRUNC` is set and the file previously existed, upon successful completion, `open()` marks for update the `st_ctime` and `st_mtime` fields of the file.

#### SYSTEM V DESCRIPTION

If *path* points to an empty string an error results.

The flags above behave as described, with the following exception:

If the `O_NDELAY` or `O_NONBLOCK` flag is set on a call to `open()`, the corresponding flag is set for that file descriptor (see `fcntl(2V)`) and subsequent reads and writes to that descriptor will not block (see `read(2V)` and `write(2V)`).

#### RETURN VALUES

`open()` returns a non-negative file descriptor on success. On failure, it returns `-1` and sets `errno` to indicate the error.

## ERRORS

EACCES	<p>Search permission is denied for a component of the path prefix of <i>path</i>.</p> <p>The file referred to by <i>path</i> does not exist, <code>O_CREAT</code> is specified, and the directory in which it is to be created does not permit writing.</p> <p><code>O_TRUNC</code> is specified and write permission is denied for the file named by <i>path</i>.</p> <p>The required permissions (for reading and/or writing) are denied for the file named by <i>path</i>.</p>
EDQUOT	<p>The file does not exist, <code>O_CREAT</code> is specified, and the directory in which the entry for the new file is being placed cannot be extended because the user's quota of disk blocks on the file system containing the directory has been exhausted.</p> <p>The file does not exist, <code>O_CREAT</code> is specified, and the user's quota of inodes on the file system on which the file is being created has been exhausted.</p>
EEXIST	<code>O_EXCL</code> and <code>O_CREAT</code> were both specified and the file exists.
EFAULT	<i>path</i> points outside the process's allocated address space.
EINTR	A signal was caught during the <code>open()</code> system call.
EIO	<p>A hangup or error occurred during a STREAMS <code>open()</code>.</p> <p>An I/O error occurred while reading from or writing to the file system.</p>
EISDIR	The named file is a directory, and the arguments specify it is to be opened for writing.
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
EMFILE	The system limit for open file descriptors per process has already been reached.
ENAMETOOLONG	<p>The length of the path argument exceeds <code>{PATH_MAX}</code>.</p> <p>A pathname component is longer than <code>{NAME_MAX}</code> while <code>{_POSIX_NO_TRUNC}</code> is in effect (see <code>pathconf(2V)</code>).</p>
ENFILE	The system file table is full.
ENOENT	<p><code>O_CREAT</code> is not set and the named file does not exist.</p> <p>A component of the path prefix of <i>path</i> does not exist.</p>
ENOSPC	<p>The file does not exist, <code>O_CREAT</code> is specified, and the directory in which the entry for the new file is being placed cannot be extended because there is no space left on the file system containing the directory.</p> <p>The file does not exist, <code>O_CREAT</code> is specified, and there are no free inodes on the file system on which the file is being created.</p>
ENOSR	A <i>stream</i> could not be allocated.
ENOTDIR	A component of the path prefix of <i>path</i> is not a directory.
ENXIO	<p><code>O_NDELAY</code> is set, the named file is a FIFO, <code>O_WRONLY</code> is set, and no process has the file open for reading.</p> <p>The file is a character special or block special file, and the associated device does not exist.</p> <p><code>O_NONBLOCK</code> is set, the named file is a FIFO, <code>O_WRONLY</code> is set, and no process has the file open for reading.</p> <p>A STREAMS module or driver open routine failed.</p>
EOPNOTSUPP	An attempt was made to open a socket (not currently implemented).

**EROFS**

The named file does not exist, **O\_CREAT** is specified, and the file system on which it is to be created is a read-only file system.

The named file resides on a read-only file system, and the file is to be opened for writing.

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

**ENOENT**                    *path* points to an empty string.

**SEE ALSO**

**chmod(2V)**, **close(2V)**, **creat(2V)**, **dup(2V)**, **fcntl(2V)**, **getdtablesize(2)**, **getmsg(2)**, **lseek(2V)**, **mknod(2V)**, **mount(2V)**, **putmsg(2)**, **read(2V)**, **umask(2V)** **write(2V)**

## NAME

pathconf, fpathconf – query file system related limits and options

## SYNOPSIS

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

```
long pathconf(path, name)
```

```
char *path;
```

```
int name;
```

```
long fpathconf(fd, name)
```

```
int fd, name;
```

## DESCRIPTION

**pathconf()** and **fpathconf()** provide a method for the application to determine the current value of a configurable limit or option that is associated with a file or directory.

For **pathconf()**, *path* points to the pathname of a file or directory. For **fpathconf()**, *fd* is an open file descriptor.

The convention used throughout sections 2 and 3 is that {LIMIT} means that LIMIT is something that can change from file to file (due to multiple file systems on the same machine). The actual value for LIMIT is typically not defined in any header file since it is not invariant. Instead, **pathconf** must be called to retrieve the value. **pathconf()** understands a list of flags that are named similarly to the value being queried.

The following table lists the name and meaning of each conceptual limit.

<i>Limit</i>	<i>Meaning</i>
{LINK_MAX}	Max links to an object.
{MAX_CANON}	Max tty input line size.
{MAX_INPUT}	Max packet a tty can accept at once.
{NAME_MAX}	Max filename length.
{PATH_MAX}	Max pathname length.
{PIPE_BUF}	Pipe buffer size.
{_POSIX_CHOWN_RESTRICTED}	If true only root can chown() files, otherwise anyone may give away files.
{_POSIX_NO_TRUNC}	If false filenames > {NAME_MAX} are truncated, otherwise an error.
{_POSIX_VDISABLE}	A char to use to disable tty special chars.

The following table lists the name of each limit, the flag passed to **pathconf()** to retrieve the value of each variable, and some notes about usage.

<i>Limit</i>	<i>Pathconf Flag</i>	<i>Notes</i>
{LINK_MAX}	_PC_LINK_MAX	1
{MAX_CANON}	_PC_MAX_CANON	2
{MAX_INPUT}	_PC_MAX_INPUT	2
{NAME_MAX}	_PC_NAME_MAX	3,4
{PATH_MAX}	_PC_PATH_MAX	4,5
{PIPE_BUF}	_PC_PIPE_BUF	6
{_POSIX_CHOWN_RESTRICTED}	_PC_CHOWN_RESTRICTED	7,8
{_POSIX_NO_TRUNC}	_PC_NO_TRUNC	3,4,8
{_POSIX_VDISABLE}	_PC_VDISABLE	2,8

The following notes apply to the entries in the preceding table.

- 1 If *path* or *fd* refers to a directory, the value returned applies to the directory itself.
- 2 The behavior is undefined if *path* or *fd* does not refer to a terminal file.
- 3 If *path* or *fd* refers to a directory, the value returned applies to the file names within the directory.

- 4 The behavior is undefined if *path* or *fd* does not refer to a directory.
- 5 If *path* or *fd* refers to a directory, the value returned is the maximum length of a relative pathname when the specified directory is the working directory.
- 6 If *path* refers to a FIFO, or *fd* refers to a pipe of FIFO, the value returned applies to the referenced object itself. If *path* or *fd* refers to a directory, the value returned applies to any FIFOs that exist or can be created within the directory. If *path* or *fd* refer to any other type of file, the behavior is undefined.
- 7 If *path* or *fd* refer to a directory, the value returned applies to any files, other than directories, that exist or can be created within the directory.
- 8 The option in question is a boolean; the return value is 0 or 1.

**RETURN VALUES**

On success, `pathconf()` and `fpathconf()` return the current variable value for the file or directory. On failure, they return `-1` and set `errno` to indicate the error.

If the variable corresponding to *name* has no limit for the path or file descriptor, `pathconf()` and `fpathconf()` return `-1` without changing `errno`.

**ERRORS**

`pathconf()` and `fpathconf()` may set `errno` to:

`EINVAL` The value of *name* is invalid.

For each of the following conditions, if the condition is detected, `pathconf()` fails and sets `errno` to:

`EACCES` Search permission is denied for a component of the path prefix.

`EINVAL` The implementation does not support an association of the variable name with the specified file.

`ENAMETOOLONG` The length of the path argument exceeds `{PATH_MAX}`.

A pathname component is longer than `{NAME_MAX}` while `{POSIX_NO_TRUNC}` is in effect.

`ENOENT` The named file does not exist.

*path* points to an empty string.

`ENOTDIR` A component of the path prefix is not a directory.

For each of the following conditions, if the condition is detected, `fpathconf()` fails and sets `errno` to:

`EBADF`

The *fd* argument is not a valid file descriptor.

`EINVAL`

The implementation does not support an association of the variable name with the specified file.

**NAME**

pipe – create an interprocess communication channel

**SYNOPSIS**

```
int pipe(fd)
int fd[2];
```

**DESCRIPTION**

The `pipe()` system call creates an I/O mechanism called a pipe and returns two file descriptors, `fd[0]` and `fd[1]`. `fd[0]` is opened for reading and `fd[1]` is opened for writing. The `O_NONBLOCK` flag is clear on both file descriptors (see `open(2V)`). When the pipe is written using the descriptor `fd[1]` up to `{PIPE_BUF}` (see `sysconf(2V)`) bytes of data are buffered before the writing process is blocked. A read only file descriptor `fd[0]` accesses the data written to `fd[1]` on a FIFO (first-in-first-out) basis.

The standard programming model is that after the pipe has been set up, two (or more) cooperating processes (created by subsequent `fork(2V)` calls) will pass data through the pipe using `read(2V)` and `write(2V)`.

Read calls on an empty pipe (no buffered data) with only one end (all write file descriptors closed) returns an EOF (end of file).

Pipes are really a special case of the `socketpair(2)` call and, in fact, are implemented as such in the system.

A `SIGPIPE` signal is generated if a write on a pipe with only one end is attempted.

Upon successful completion, `pipe()` marks for update the `st_atime`, `st_ctime`, and `st_mtime` fields of the pipe.

**RETURN VALUES**

`pipe()` returns:

- 0        on success.
- 1       on failure and sets `errno` to indicate the error.

**ERRORS**

- EFAULT        The array `fd` is in an invalid area of the process's address space.
- EMFILE        Too many descriptors are active.
- ENFILE        The system file table is full.

**SEE ALSO**

`sh(1)`, `fork(2V)`, `read(2V)`, `socketpair(2)`, `write(2V)`

**BUGS**

Should more than `{PIPE_BUF}` bytes be necessary in any pipe among a loop of processes, deadlock will occur.



## NAME

`poll` – I/O multiplexing

## SYNOPSIS

```
#include <poll.h>

int poll(fds, nfd, timeout)
struct pollfd *fds;
unsigned long nfd;
int timeout;
```

## DESCRIPTION

`poll()` provides users with a mechanism for multiplexing input/output over a set of file descriptors (see `intro(2)`). `poll()` identifies those file descriptors on which a user can send or receive messages, or on which certain events have occurred. A user can receive messages using `read(2V)` or `getmsg(2)` and can send messages using `write(2V)` and `putmsg(2)`. Certain `ioctl(2)` calls, such as `I_RECVFD` and `I_SENDFD` (see `streamio(4)`), can also be used to receive and send messages on streams.

`fds` specifies the file descriptors to be examined and the events of interest for each file descriptor. It is a pointer to an array with one element for each open file descriptor of interest. The array's elements are `pollfd` structures which contain the following members:

```
int fd;           /* file descriptor */
short events;    /* requested events */
short revents;   /* returned events */
```

where `fd` specifies an open file descriptor and `events` and `revents` are bitmasks constructed by ORing any combination of the following event flags:

- POLLIN** If the file descriptor refers to a stream, a non-priority or file descriptor passing message (see `I_RECVFD`) is present on the **stream head** read queue. This flag is set even if the message is of zero length. If the file descriptor is not a stream, the file descriptor is readable. In `revents`, this flag is mutually exclusive with **POLLPRI**.
- POLLPRI** If the file descriptor is a stream, a priority message is present on the **stream head** read queue. This flag is set even if the message is of zero length. If the file descriptor is not a stream, some exceptional condition has occurred. In `revents`, this flag is mutually exclusive with **POLLIN**.
- POLLOUT** If the file descriptor is a stream, the first downstream write queue in the *stream* is not full. Priority control messages can be sent (see `putmsg(2)`) at any time. If the file descriptor is not a stream, it is writable.
- POLLERR** If the file descriptor is a stream, an error message has arrived at the **stream head**. This flag is only valid in the `revents` bitmask; it is not used in the `events` field.
- POLLHUP** If the file descriptor is a stream, a hangup has occurred on the *stream*. This event and **POLLOUT** are mutually exclusive; a *stream* can never be writable if a hangup has occurred. However, this event and **POLLIN** or **POLLPRI** are not mutually exclusive. This flag is only valid in the `revents` bitmask; it is not used in the `events` field.
- POLLNVAL** The specified `fd` value does not specify an open file descriptor. This flag is only valid in the `revents` field; it is not used in the `events` field.

For each element of the array pointed to by `fds`, `poll()` examines the given file descriptor for the `event(s)` specified in `events`. The number of file descriptors to be examined is specified by `nfd`. If `nfd` exceeds the system limit of open files (see `getdtablesize(2)`), `poll()` will fail.

If the value `fd` is less than zero, `events` is ignored and `revents` is set to 0 in that entry on return from `poll()`.

The results of the `poll()` query are stored in the `revents` field in the `pollfd` structure. Bits are set in the `revents` bitmask to indicate which of the requested events are true. If none are true, none of the specified bits is set in `revents` when the `poll()` call returns. The event flags `POLLHUP`, `POLLERR`, and `POLLNVAL` are always set in `revents` if the conditions they indicate are true; this occurs even though these flags were not present in `events`.

If none of the defined events have occurred on any selected file descriptor, `poll()` waits at least *timeout* milliseconds for an event to occur on any of the selected file descriptors. On a computer where millisecond timing accuracy is not available, *timeout* is rounded up to the nearest legal value available on that system. If the value *timeout* is 0, `poll()` returns immediately. If the value of *timeout* is -1, `poll()` blocks until a requested event occurs or until the call is interrupted. `poll()` is not affected by the `O_NDELAY` flag.

#### RETURN VALUES

`poll()` returns a non-negative value on success. A positive value indicates the total number of file descriptors that has been selected (for instance, file descriptors for which the `revents` field is non-zero). 0 indicates the call timed out and no file descriptors have been selected. On failure, `poll()` returns -1 and sets `errno` to indicate the error.

#### ERRORS

EAGAIN	Allocation of internal data structures failed, but the request should be attempted again.
EFAULT	Some argument points outside the allocated address space.
EINTR	A signal was caught during the <code>poll()</code> system call.
EINVAL	The argument <i>nfds</i> is less than zero. <i>nfds</i> is greater than the system limit of open files.

#### SEE ALSO

`getdtablesize(2)`, `getmsg(2)`, `intro(2)`, `ioctl(2)`, `putmsg(2)`, `read(2V)`, `select(2)`, `write(2V)`, `streamio(4)`

**NAME**

profil – execution time profile

**SYNOPSIS**

```
int profil(buf, bufsiz, offset, scale)
short *buf;
int bufsiz;
void (*offset)();
int scale;
```

**DESCRIPTION**

**profil()** enables run-time execution profiling, and reserves a buffer for maintaining raw profiling statistics. *buf* points to an area of core of length *bufsiz* (in bytes). After the call to **profil()**, the user's program counter (*pc*) is examined at each clock tick (10 milliseconds on Sun-4 systems, 20 milliseconds on Sun-3 systems); *offset* is subtracted from its value, and the result multiplied by *scale*. If the resulting number corresponds to a word within the buffer, that word is incremented.

*scale* is interpreted as an unsigned, fixed-point fraction with binary point at the left: 0xffff gives a 1-to-1 mapping of *pc* values to words in *buf*; 0x7fff maps each pair of instruction words together. 0x2 maps all instructions onto the beginning of *buf* (producing a non-interrupting core clock).

Profiling is turned off by giving a *scale* of 0 or 1. It is rendered ineffective by giving a *bufsiz* of 0. Profiling is turned off when an **execve()** is executed, but remains on in child and parent both after a **fork()**. Profiling is turned off if an update in *buf* would cause a memory fault.

**RETURN VALUES**

**profil()** always succeeds and returns 0.

**SEE ALSO**

**gprof(1)**, **getitimer(2)**, **monitor(3)**

## NAME

ptrace – process trace

## SYNOPSIS

```
#include <signal.h>
#include <sys/ptrace.h>
#include <sys/wait.h>

ptrace(request, pid, addr, data [ , addr2 ] )
enum ptracereq request;
int pid;
char *addr;
int data;
char *addr2;
```

## DESCRIPTION

**ptrace()** provides a means by which a process may control the execution of another process, and examine and change its core image. Its primary use is for the implementation of breakpoint debugging. There are five arguments whose interpretation depends on the *request* argument. Generally, *pid* is the process ID of the traced process. A process being traced behaves normally until it encounters some signal whether internally generated like “illegal instruction” or externally generated like “interrupt”. See **sigvec(2)** for the list. Then the traced process enters a stopped state and the tracing process is notified using **wait(2V)**. When the traced process is in the stopped state, its core image can be examined and modified using **ptrace()**. If desired, another **ptrace()** request can then cause the traced process either to terminate or to continue, possibly ignoring the signal.

Note: several different values of the *request* argument can make **ptrace()** return data values — since `-1` is a possibly legitimate value, to differentiate between `-1` as a legitimate value and `-1` as an error code, you should clear the **errno** global error code before doing a **ptrace()** call, and then check the value of **errno** afterwards.

The value of the *request* argument determines the precise action of the call:

**PTRACE\_TRACEME**

This request is the only one used by the traced process; it declares that the process is to be traced by its parent. All the other arguments are ignored. Peculiar results will ensue if the parent does not expect to trace the child.

**PTRACE\_PEEKTEXT****PTRACE\_PEEKDATA**

The word in the traced process’s address space at *addr* is returned. If the instruction and data spaces are separate (for example, historically on a PDP-11), request **PTRACE\_PEEKTEXT** indicates instruction space while **PTRACE\_PEEKDATA** indicates data space. Otherwise, either request may be used, with equal results; *addr* must be a multiple of 4 on a Sun-4 system. The child must be stopped. The input *data* and *addr2* are ignored.

**PTRACE\_PEEKUSER**

The word of the system’s per-process data area corresponding to *addr* is returned. *addr* must be a valid offset within the kernel’s per-process data pages. This space contains the registers and other information about the process; its layout corresponds to the *user* structure in the system (see **<sys/user.h>**).

**PTRACE\_POKETEXT****PTRACE\_POKEDATA**

The given *data* are written at the word in the process’s address space corresponding to *addr*. *addr* must be a multiple of 4 on a Sun-4 system. No useful value is returned. If the instruction and data spaces are separate, request **PTRACE\_PEEKTEXT** indicates instruction space while **PTRACE\_PEEKDATA** indicates data space. The **PTRACE\_POKETEXT** request must be used to write into a process’s text space even if the instruction and data spaces are not separate.

**PTRACE\_POKEUSER**

The process's system data are written, as it is read with request **PTRACE\_PEEKUSER**. Only a few locations can be written in this way: the general registers, the floating point and status registers, and certain bits of the processor status word.

**PTRACE\_CONT**

The *data* argument is taken as a signal number and the child's execution continues at location *addr* as if it had incurred that signal. Normally the signal number will be either 0 to indicate that the signal that caused the stop should be ignored, or that value fetched out of the process's image indicating which signal caused the stop. If *addr* is (int \*)1 then execution continues from where it stopped. *addr* must be a multiple of 4 on a Sun-4 system.

**PTRACE\_KILL**

The traced process terminates, with the same consequences as **exit(2V)**.

**PTRACE\_SINGLESTEP**

Execution continues as in request **PTRACE\_CONT**; however, as soon as possible after execution of at least one instruction, execution stops again. The signal number from the stop is **SIGTRAP**. On Sun-3 and Sun386i systems, the status register T-bit is used and just one instruction is executed. This is part of the mechanism for implementing breakpoints. On a Sun-4 system this will return an error since there is no hardware assist for this feature. Instead, the user should insert breakpoint traps in the debugged program with **PTRACE\_POKETEXT**.

**PTRACE\_ATTACH**

Attach to the process identified by the *pid* argument and begin tracing it. **PTRACE\_ATTACH** causes a **SIGSTOP** to be sent to process *pid*. Process *pid* does not have to be a child of the requestor, but the requestor must have permission to send process *pid* a signal and the effective user IDs of the requesting process and process *pid* must match.

**PTRACE\_DETACH**

Detach the process being traced. Process *pid* is no longer being traced and continues its execution. The *data* argument is taken as a signal number and the process continues at location *addr* as if it had incurred that signal.

**PTRACE\_GETREGS**

The traced process's registers are returned in a structure pointed to by the *addr* argument. The registers include the general purpose registers, the program counter and the program status word. The "regs" structure defined in <machine/reg.h> describes the data that are returned.

**PTRACE\_SETREGS**

The traced process's registers are written from a structure pointed to by the *addr* argument. The registers include the general purpose registers, the program counter and the program status word. The "regs" structure defined in **reg.h** describes the data that are set.

**PTRACE\_GETFPREGS**

(Sun-3, Sun-4 and Sun386i systems only) The traced process's FPP status is returned in a structure pointed to by the *addr* argument. The status includes the 68881 (80387 on Sun386i systems) floating point registers and the control, status, and instruction address registers. The "fp\_status" structure defined in **reg.h** describes the data that are returned. The **fp\_state** structure defined in <machine/fp.h> describes the data that are returned on a Sun386i system.

**PTRACE\_SETFPREGS**

(Sun-3, Sun-4 and Sun386i systems only) The traced process's FPP status is written from a structure pointed to by the *addr* argument. The status includes the FPP floating point registers and the control, status, and instruction address registers. The "fp\_status" structure defined in **reg.h** describes the data that are set. The "fp\_state" structure defined in **fp.h** describes the data that are returned on a Sun386i system.

**PTRACE\_GETFPAREGS**

(a Sun-3 system with FPA only) The traced process's FPA registers are returned in a structure pointed to by the *addr* argument. The "fpa\_regs" structure defined in *reg.h* describes the data that are returned.

**PTRACE\_SETFPAREGS**

(a Sun-3 system with FPA only) The traced process's FPA registers are written from a structure pointed to by the *addr* argument. The "fpa\_regs" structure defined in *reg.h* describes the data that are set.

**PTRACE\_READTEXT****PTRACE\_READDATA**

Read data from the address space of the traced process. If the instruction and data spaces are separate, request **PTRACE\_READTEXT** indicates instruction space while **PTRACE\_READDATA** indicates data space. The *addr* argument is the address within the traced process from where the data are read, the *data* argument is the number of bytes to read, and the *addr2* argument is the address within the requesting process where the data are written.

**PTRACE\_WRITETEXT****PTRACE\_WRITEDATA**

Write data into the address space of the traced process. If the instruction and data spaces are separate, request **PTRACE\_READTEXT** indicates instruction space while **PTRACE\_READDATA** indicates data space. The *addr* argument is the address within the traced process where the data are written, the *data* argument is the number of bytes to write, and the *addr2* argument is the address within the requesting process from where the data are read.

**PTRACE\_SETWRBKPT**

(Sun386i systems only) Set a write breakpoint at location *addr* in the process being traced. Whenever a write is directed to this location a breakpoint will occur and a SIGTRAP signal will be sent to the process. The *data* argument specifies which debug register should be used for the address of the breakpoint and must be in the range 0 through 3, inclusive. The *addr2* argument specifies the length of the operand in bytes, and must be one of 1, 2, or 4.

**PTRACE\_SETACBKPT**

(Sun386i systems only) Set an access breakpoint at location *addr* in the process being traced. When location *addr* is read or written a breakpoint will occur and the process will be sent a SIGTRAP signal. The *data* argument specifies which debug register should be used for the address of the breakpoint and must be in the range 0 through 3, inclusive. The *addr2* argument specifies the length of the operand in bytes, and must be one of 1, 2, or 4.

**PTRACE\_CLRBKPT**

(Sun386i systems only) Clears all break points set with **PTRACE\_SETACBKPT** or **PTRACE\_SETWRBKPT**.

**PTRACE\_SYSCALL**

Execution continues as in request **PTRACE\_CONT**; until the process makes a system call. The process receives a SIGTRAP signal and stops. At this point the arguments to the system call may be inspected in the process *user* structure using the **PTRACE\_PEEKUSER** request. The system call number is available in place of the 8th argument. Continuing with another **PTRACE\_SYSCALL** will stop the process again at the completion of the system call. At this point the result of the system call and error value may be inspected in the process *user* structure.

**PTRACE\_DUMPCORE**

Dumps a core image of the traced process to a file. The name of the file is obtained from the *addr* argument.

As indicated, these calls (except for requests `PTRACE_TRACEME`, `PTRACE_ATTACH` and `PTRACE_DETACH`) can be used only when the subject process has stopped. The `wait()` call is used to determine when a process stops; in such a case the “termination” status returned by `wait()` has the value `WSTOPPED` to indicate a stop rather than genuine termination.

To forestall possible fraud, `ptrace()` inhibits the `setUID` and `setGID` facilities on subsequent `execve(2V)` calls. If a traced process calls `execve()`, it will stop before executing the first instruction of the new image, showing signal `SIGTRAP`.

On the Sun, “word” also means a 32-bit integer.

#### RETURN VALUES

On success, the value returned by `ptrace()` depends on *request* as follows:

<code>PTRACE_PEEKTEXT</code>	
<code>PTRACE_PEEKDATA</code>	The word in the traced process’s address space at <i>addr</i> .
<code>PTRACE_PEEKUSER</code>	The word of the system’s per-process data area corresponding to <i>addr</i> .

On failure, these requests return `-1` and set `errno` to indicate the error.

For all other values of *request*, `ptrace()` returns:

0	on success.
-1	on failure and sets <code>errno</code> to indicate the error.

#### ERRORS

<code>EIO</code>	The request code is invalid.
	The given signal number is invalid.
	The specified address is out of bounds.
<code>EPERM</code>	The specified process cannot be traced.
<code>ESRCH</code>	The specified process does not exist.
	<i>request</i> requires process <i>pid</i> to be traced by the current process and stopped, and process <i>pid</i> is not being traced by the current process.
	<i>request</i> requires process <i>pid</i> to be traced by the current process and stopped, and process <i>pid</i> is not stopped.

#### SEE ALSO

`adb(1)`, `intro(2)`, `ioctl(2)`, `sigvec(2)`, `wait(2V)`

#### BUGS

`ptrace()` is unique and arcane; it should be replaced with a special file which can be opened and read and written. The control functions could then be implemented with `ioctl(2)` calls on this file. This would be simpler to understand and have much higher performance.

The requests `PTRACE_TRACEME` through `PTRACE_SINGLESTEP` are standard UNIX system `ptrace()` requests. The requests `PTRACE_ATTACH` through `PTRACE_DUMPCORE` and the fifth argument, *addr2*, are unique to SunOS.

The request `PTRACE_TRACEME` should be able to specify signals which are to be treated normally and not cause a stop. In this way, for example, programs with simulated floating point (which use “illegal instruction” signals at a very high rate) could be efficiently debugged.

The error indication, `-1`, is a legitimate function value; `errno`, (see `intro(2)`), can be used to clarify what it means.

## NAME

putmsg – send a message on a stream

## SYNOPSIS

```
#include <stropts.h>

int putmsg(fd, ctlptr, dataptr, flags)
int fd;
struct strbuf *ctlptr;
struct strbuf *dataptr;
int flags;
```

## DESCRIPTION

**putmsg()** creates a message (see **intro(2)**) from user specified buffer(s) and sends the message to a STREAMS file. The message may contain either a data part, a control part or both. The data and control parts to be sent are distinguished by placement in separate buffers, as described below. The semantics of each part is defined by the STREAMS module that receives the message.

*fd* specifies a file descriptor referencing an open *stream*. *ctlptr* and *dataptr* each point to a **strbuf** structure that contains the following members:

```
int maxlen;    /* not used */
int len;       /* length of data */
char *buf;     /* ptr to buffer */
```

*ctlptr* points to the structure describing the control part, if any, to be included in the message. The **buf** field in the **strbuf** structure points to the buffer where the control information resides, and the **len** field indicates the number of bytes to be sent. The **maxlen** field is not used in **putmsg()** (see **getmsg(2)**). In a similar manner, *dataptr* specifies the data, if any, to be included in the message. *flags* may be set to the values 0 or **RS\_HIPRI** and is used as described below.

To send the data part of a message, *dataptr* must not be a NULL pointer and the **len** field of *dataptr* must have a value of 0 or greater. To send the control part of a message, the corresponding values must be set for *ctlptr*. No data (control) part will be sent if either *dataptr* (*ctlptr*) is a NULL pointer or the **len** field of *dataptr* (*ctlptr*) is set to -1.

If a control part is specified, and *flags* is set to **RS\_HIPRI**, a *priority* message is sent. If *flags* is set to 0, a non-priority message is sent. If no control part is specified, and *flags* is set to **RS\_HIPRI**, **putmsg()** fails and sets **errno** to **EINVAL**. If no control part and no data part are specified, and *flags* is set to 0, no message is sent, and 0 is returned.

For non-priority messages, **putmsg()** will block if the *stream* write queue is full due to internal flow control conditions. For priority messages, **putmsg()** does not block on this condition. For non-priority messages, **putmsg()** does not block when the write queue is full and **O\_NDELAY** is set. Instead, it fails and sets **errno** to **EAGAIN**.

**putmsg()** also blocks, unless prevented by lack of internal resources, waiting for the availability of message blocks in the *stream*, regardless of priority or whether **O\_NDELAY** has been specified. No partial message is sent.

## RETURN VALUES

**putmsg()** returns:

- 0        on success.
- 1       on failure and sets **errno** to indicate the error.

## ERRORS

- EAGAIN**        A non-priority message was specified, the **O\_NDELAY** flag is set and the *stream* write queue is full due to internal flow control conditions.  
Buffers could not be allocated for the message that was to be created.



EBADF	<i>fd</i> is not a valid file descriptor open for writing.
EFAULT	<i>ctlptr</i> or <i>dataptr</i> points outside the allocated address space.
EINTR	A signal was caught during the <b>putmsg()</b> system call.
EINVAL	An undefined value was specified in <i>flags</i> . <i>flags</i> is set to <b>RS_HIPRI</b> and no control part was supplied. The <i>stream</i> referenced by <i>fd</i> is linked below a multiplexor.
ENOSTR	A <i>stream</i> is not associated with <i>fd</i> .
ENXIO	A hangup condition was generated downstream for the specified <i>stream</i> .
ERANGE	The size of the data part of the message does not fall within the range specified by the maximum and minimum packet sizes of the topmost <i>stream</i> module. The control part of the message is larger than the maximum configured size of the control part of a message. The data part of the message is larger than the maximum configured size of the data part of a message.

A **putmsg()** also fails if a STREAMS error message had been processed by the *stream* head before the call to **putmsg()**. The error returned is the value contained in the STREAMS error message.

**SEE ALSO**

**getmsg(2), intro(2), poll(2), read(2V), write(2V)**

**NAME**

quotactl – manipulate disk quotas

**SYNOPSIS**

```
#include <ufs/quota.h>
```

```
int quotactl(cmd, special, uid, addr)
```

```
int cmd;
```

```
char *special;
```

```
int uid;
```

```
caddr_t addr;
```

**DESCRIPTION**

The **quotactl()** call manipulates disk quotas. *cmd* indicates a command to be applied to the user ID *uid*. *special* is a pointer to a null-terminated string containing the path name of the block special device for the file system being manipulated. The block special device must be mounted as a UFS file system (see **mount(2V)**). *addr* is the address of an optional, command specific, data structure which is copied in or out of the system. The interpretation of *addr* is given with each command below.

- Q\_QUOTAON** Turn on quotas for a file system. *addr* points to the path name of file containing the quotas for the file system. The quota file must exist; it is normally created with the **quota-check(8)** program. This call is restricted to the super-user.
- Q\_QUOTAOFF** Turn off quotas for a file system. *addr* and *uid* are ignored. This call is restricted to the super-user.
- Q\_GETQUOTA** Get disk quota limits and current usage for user *uid*. *addr* is a pointer to a **dqblk** structure (defined in **<ufs/quota.h>**). Only the super-user may get the quotas of a user other than himself.
- Q\_SETQUOTA** Set disk quota limits and current usage for user *uid*. *addr* is a pointer to a **dqblk** structure (defined in **quota.h**). This call is restricted to the super-user.
- Q\_SETQLIM** Set disk quota limits for user *uid*. *addr* is a pointer to a **dqblk** structure (defined in **quota.h**). This call is restricted to the super-user.
- Q\_SYNC** Update the on-disk copy of quota usages for a file system. If *special* is null then all file systems with active quotas are sync'ed. *addr* and *uid* are ignored.

**RETURN VALUES**

**quotactl()** returns:

- 0 on success.
- 1 on failure and sets **errno** to indicate the error.

**ERRORS**

- EFAULT** *addr* or *special* are invalid.
- EINVAL** The kernel has not been compiled with the QUOTA option.  
*cmd* is invalid.
- ENODEV** *special* is not a mounted UFS file system.
- ENOENT** The file specified by *special* or *addr* does not exist.
- ENOTBLK** *special* is not a block device.
- EPERM** The call is privileged and the caller was not the super-user.
- ESRCH** No disc quota is found for the indicated user.  
Quotas have not been turned on for this file system.
- EUSERS** The quota table is full.

If *cmd* is `Q_QUOTAON` `quotactl()` may set `errno` to:

- `EACCES`        The quota file pointed to by *addr* exists but is not a regular file.  
                  The quota file pointed to by *addr* exists but is not on the file system pointed to by *special*.
- `EBUSY`        `Q_QUOTAON` attempted while another `Q_QUOTAON` or `Q_QUOTAOFF` is in progress.

**SEE ALSO**

`quota(1)`, `getrlimit(2)`, `mount(2V)`, `quotacheck(8)`, `quotaon(8)`

**BUGS**

There should be some way to integrate this call with the resource limit interface provided by `setrlimit()` and `getrlimit(2)`.

Incompatible with Melbourne quotas.

## NAME

read, readv – read input

## SYNOPSIS

```
int read(fd, buf, nbyte)
```

```
int fd;
```

```
char *buf;
```

```
int nbyte;
```

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

```
#include <sys/uio.h>
```

```
int readv(fd, iov, iovcnt)
```

```
int fd;
```

```
struct iovec *iov;
```

```
int iovcnt;
```

## DESCRIPTION

**read()** attempts to read *nbyte* bytes of data from the object referenced by the descriptor *fd* into the buffer pointed to by *buf*. **readv()** performs the same action as **read()**, but scatters the input data into the *iovcnt* buffers specified by the members of the *iov* array: *iov*[0], *iov*[1], ..., *iov*[*iovcnt* - 1].

If *nbyte* is zero, **read()** takes no action and returns 0. **readv()**, however, returns -1 and sets the global variable *errno* (see ERRORS below).

For **readv()**, the *iovec* structure is defined as

```
struct iovec {
    caddr_t iov_base;
    int     iov_len;
};
```

Each *iovec* entry specifies the base address and length of an area in memory where data should be placed. **readv()** will always fill an area completely before proceeding to the next.

On objects capable of seeking, the **read()** starts at a position given by the pointer associated with *fd* (see **lseek(2V)**). Upon return from **read()**, the pointer is incremented by the number of bytes actually read.

Objects that are not capable of seeking always read from the current position. The value of the pointer associated with such an object is undefined.

Upon successful completion, **read()** and **readv()** return the number of bytes actually read and placed in the buffer. The system guarantees to read the number of bytes requested if the descriptor references a normal file which has that many bytes left before the EOF (end of file), but in no other case.

If the process calling **read()** or **readv()** receives a signal before any data are read, the system call is restarted unless the process explicitly set the signal to interrupt the call using **sigvec()** or **sigaction()** (see the discussions of **SV\_INTERRUPT** on **sigvec(2)** and **SA\_INTERRUPT** on **sigaction(3V)**). If **read()** or **readv()** is interrupted by a signal after successfully reading some data, it returns the number of bytes read.

If *nbyte* is not zero and **read()** returns 0, then EOF has been reached. If **readv()** returns 0, then EOF has been reached.

A **read()** or **readv()** from a STREAMS file (see **intro(2)**) can operate in three different modes: “byte-stream” mode, “message-nondiscard” mode, and “message-discard” mode. The default is byte-stream mode. This can be changed using the **I\_SRDOPT ioctl(2)** request (see **streamio(4)**), and can be tested with the **I\_GRDOPT ioctl()** request. In byte-stream mode, **read()** and **readv()** will retrieve data from the *stream* until as many bytes as were requested are transferred, or until there is no more data to be retrieved. Byte-stream mode ignores message boundaries.

In STREAMS message-nondiscard mode, **read()** and **readv()** will retrieve data until as many bytes as were requested are transferred, or until a message boundary is reached. If the **read()** or **readv()** does not retrieve all the data in a message, the remaining data are left on the *stream*, and can be retrieved by the

next `read()`, `readv()`, or `getmsg(2)` call. Message-discard mode also retrieves data until as many bytes as were requested are transferred, or a message boundary is reached. However, unread data remaining in a message after the `read()` or `readv()` returns are discarded, and are not available for a subsequent `read()`, `readv()`, or `getmsg()`.

When attempting to read from a descriptor associated with an empty pipe, socket, FIFO, or *stream*:

- If the object the descriptor is associated with is marked for 4.2BSD-style non-blocking I/O (with the `FIONBIO ioctl()` request or a call to `fcntl(2V)` using the `FNDELAY` flag from `<sys/file.h>` or the `O_NDELAY` flag from `<fcntl.h>` in the 4.2BSD environment), the read will return `-1` and `errno` will be set to `EWOULDBLOCK`.
- If the descriptor is marked for System V-style non-blocking I/O (using `fcntl()` with the `FNDELAY` flag from `<sys/file.h>` or the `O_NDELAY` flag from `<fcntl.h>` in the System V environment), and does not refer to a *stream*, the read will return `0`. Note: this is indistinguishable from EOF.
- If the descriptor is marked for POSIX-style non-blocking I/O (using `fcntl()` with the `O_NONBLOCK` flag from `<fcntl.h>`) and refers to a *stream*, the read will return `-1` and `errno` will be set to `EAGAIN`.
- If neither the descriptor nor the object it refers to are marked for non-blocking I/O, the read will block until data is available to be read or the object has been “disconnected”. A pipe or FIFO is “disconnected” when no process has the object open for writing; a socket that was connected is “disconnected” when the connection is broken; a *stream* is “disconnected” when a hangup condition occurs (for instance, when carrier drops on a terminal).

If the descriptor or the object is marked for non-blocking I/O, and less data are available than are requested by the `read()` or `readv()`, only the data that are available are returned, and the count indicates how many bytes of data were actually read.

When reading from a STREAMS file, handling of zero-byte messages is determined by the current read mode setting. In byte-stream mode, `read()` and `readv()` accept data until as many bytes as were requested are transferred, or until there is no more data to read, or until a zero-byte message block is encountered. `read()` and `readv()` then return the number of bytes read, and places the zero-byte message back on the *stream* to be retrieved by the next `read()`, `readv()`, or `getmsg()`. In the two other modes, a zero-byte message returns a value of `0` and the message is removed from the *stream*. When a zero-byte message is read as the first message on a *stream*, a value of `0` is returned regardless of the read mode.

A `read()` or `readv()` from a STREAMS file can only process data messages. It cannot process any type of protocol message and will fail if a protocol message is encountered at the *streamhead*.

Upon successful completion, `read()` and `readv()` mark for update the `st_atime` field of the file.

#### RETURN VALUES

`read()` and `readv()` return the number of bytes actually read on success. On failure, they return `-1` and set `errno` to indicate the error.

#### ERRORS

EAGAIN	The descriptor referred to a <i>stream</i> , was marked for System V-style non-blocking I/O, and no data were ready to be read.
EBADF	<i>d</i> is not a valid file descriptor open for reading.
EBADMSG	The message waiting to be read on a <i>stream</i> is not a data message.
EFAULT	<i>buf</i> points outside the allocated address space.
EINTR	The process performing a read from a slow device received a signal before any data arrived, and the signal was set to interrupt the system call.
EINVAL	The <i>stream</i> is linked below a multiplexor. The pointer associated with <i>fd</i> was negative.

- EIO** An I/O error occurred while reading from or writing to the file system.  
The calling process is in a background process group and is attempting to read from its controlling terminal and the process is ignoring or blocking SIGTTIN.  
The calling process is in a background process group and is attempting to read from its controlling terminal and the process is orphaned.
- EISDIR** *fd* refers to a directory which is on a file system mounted using the NFS.
- EWOULDBLOCK** The file was marked for 4.2BSD-style non-blocking I/O, and no data were ready to be read.

In addition to the above, `readv()` may set `errno` to:

- EFAULT** Part of *iov* points outside the process's allocated address space.
- EINVAL** *iovcnt* was less than or equal to 0, or greater than 16.  
One of the `iov_len` values in the *iov* array was negative.  
The sum of the `iov_len` values in the *iov* array overflowed a 32-bit integer.

A `read()` or `readv()` from a STREAMS file will also fail if an error message is received at the *stream* head. In this case, `errno` is set to the value returned in the error message. If a hangup occurs on the *stream* being read, `read()` will continue to operate normally until the **stream head** read queue is empty. Thereafter, it will return 0.

**SEE ALSO**

`dup(2V)`, `fcntl(2V)`, `getmsg(2)`, `intro(2)`, `ioctl(2)`, `lseek(2V)`, `open(2V)`, `pipe(2V)`, `select(2)`, `socket(2)`, `socketpair(2)`, `streamio(4)`, `termio(4)`

**NAME**

**readlink** – read value of a symbolic link

**SYNOPSIS**

```
int readlink(path, buf, bufsiz)  
char *path, *buf;  
int bufsiz;
```

**DESCRIPTION**

**readlink()** places the contents of the symbolic link referred to by *path* in the buffer *buf* which has size *bufsiz*. The contents of the link are not null terminated when returned.

**RETURN VALUES**

**readlink()** returns the number of characters placed in the buffer on success. On failure, it returns `-1` and sets `errno` to indicate the error.

**ERRORS**

**readlink()** will fail and the buffer will be unchanged if:

<b>EACCES</b>	Search permission is denied for a component of the path prefix of <i>path</i> .
<b>EFAULT</b>	<i>path</i> or <i>buf</i> extends outside the process's allocated address space.
<b>ELOOP</b>	Too many symbolic links were encountered in translating <i>path</i> .
<b>EINVAL</b>	The named file is not a symbolic link.
<b>EIO</b>	An I/O error occurred while reading from or writing to the file system.
<b>ENAMETOOLONG</b>	The length of the path argument exceeds <code>{PATH_MAX}</code> . A pathname component is longer than <code>{NAME_MAX}</code> (see <code>sysconf(2V)</code> ) while <code>{_POSIX_NO_TRUNC}</code> is in effect (see <code>pathconf(2V)</code> ).
<b>ENOENT</b>	The named file does not exist.

**SEE ALSO**

`stat(2V)`, `symlink(2)`

**NAME**

reboot – reboot system or halt processor

**SYNOPSIS**

```
#include <sys/reboot.h>

reboot(howto, [ bootargs ] )
int howto;
char *bootargs;
```

**DESCRIPTION**

**reboot()** reboots the system, and is invoked automatically in the event of unrecoverable system failures. *howto* is a mask of options passed to the bootstrap program. The system call interface permits only **RB\_HALT** or **RB\_AUTOBOOT** to be passed to the reboot program; the other flags are used in scripts stored on the console storage media, or used in manual bootstrap procedures. When none of these options (for instance **RB\_AUTOBOOT**) is given, the system is rebooted from file **/vmunix** in the root file system of unit 0 of a disk chosen in a processor specific way. An automatic consistency check of the disks is then normally performed.

The bits of *howto* are:

- |                   |  |
|-------------------|--|
| <b>RB_HALT</b>    | the processor is simply halted; no reboot takes place. <b>RB_HALT</b> should be used with caution.   |
| <b>RB_ASKNAME</b> | Interpreted by the bootstrap program itself, causing it to inquire as to what file should be booted. Normally, the system is booted from the file <b>/vmunix</b> without asking.   |
| <b>RB_SINGLE</b>  | Normally, the reboot procedure involves an automatic disk consistency check and then multi-user operations. <b>RB_SINGLE</b> prevents the consistency check, rather simply booting the system with a single-user shell on the console. <b>RB_SINGLE</b> is interpreted by the <b>init(8)</b> program in the newly booted system. |
| <b>RB_DUMP</b>    | A system core dump is performed before rebooting.  |
| <b>RB_STRING</b>  | The optional argument <i>bootargs</i> is passed to the bootstrap program. See <b>boot(8S)</b> for details. This option overrides <b>RB_SINGLE</b> but the same effect can be achieved by including <b>-s</b> as an option in <i>bootargs</i> .   |

Only the super-user may **reboot()** a machine.

**RETURN VALUES**

On success, **reboot()** does not return. On failure, it returns **-1** and sets **errno** to indicate the error.

**ERRORS**

- |              |                                   |
|--------------|-----------------------------------|
| <b>EPERM</b> | The caller is not the super-user. |
|--------------|-----------------------------------|

**FILES**

**/vmunix**

**SEE ALSO**

**panic(8S)**, **halt(8)**, **init(8)**, **intro(8)**, **reboot(8)**



## NAME

recv, recvfrom, recvmsg – receive a message from a socket

## SYNOPSIS

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

int recv(s, buf, len, flags)
int s;
char *buf;
int len, flags;

int recvfrom(s, buf, len, flags, from, fromlen)
int s;
char *buf;
int len, flags;
struct sockaddr *from;
int *fromlen;

int recvmsg(s, msg, flags)
int s;
struct msghdr *msg;
int flags;
```

## DESCRIPTION

*s* is a socket created with `socket(2)`. `recv()`, `recvfrom()`, and `recvmsg()` are used to receive messages from another socket. `recv()` may be used only on a *connected* socket (see `connect(2)`), while `recvfrom()` and `recvmsg()` may be used to receive data on a socket whether it is in a connected state or not.

If *from* is not a NULL pointer, the source address of the message is filled in. *fromlen* is a value-result parameter, initialized to the size of the buffer associated with *from*, and modified on return to indicate the actual size of the address stored there. The length of the message is returned. If a message is too long to fit in the supplied buffer, excess bytes may be discarded depending on the type of socket the message is received from (see `socket(2)`).

If no messages are available at the socket, the receive call waits for a message to arrive, unless the socket is non-blocking (see `ioctl(2)`) in which case `-1` is returned with the external variable `errno` set to `EWOULDBLOCK`.

The `select(2)` call may be used to determine when more data arrive.

If the process calling `recv()`, `recvfrom()` or `recvmsg()` receives a signal before any data are available, the system call is restarted unless the calling process explicitly set the signal to interrupt these calls using `sigvec()` or `sigaction()` (see the discussions of `SV_INTERRUPT` on `sigvec(2)`, and `SA_INTERRUPT` on `sigaction(3V)`).

The *flags* parameter is formed by ORing one or more of the following:

<code>MSG_OOB</code>	Read any “out-of-band” data present on the socket, rather than the regular “in-band” data.
<code>MSG_PEEK</code>	“Peek” at the data present on the socket; the data are returned, but not consumed, so that a subsequent receive operation will see the same data.

The `recvmsg()` call uses a `msg_hdr` structure to minimize the number of directly supplied parameters. This structure is defined in `<sys/socket.h>`, and includes the following members:

```

caddr_t  msg_name;      /* optional address */
int      msg_namelen;   /* size of address */
struct iovec *msg_iov;  /* scatter/gather array */
int      msg_iovlen;    /* # elements in msg_iov */
caddr_t  msg_accrights; /* access rights sent/received */
int      msg_accrightslen;

```

Here `msg_name` and `msg_namelen` specify the destination address if the socket is unconnected; `msg_name` may be given as a NULL pointer if no names are desired or required. The `msg_iov` and `msg_iovlen` describe the scatter-gather locations, as described in `read(2V)`. A buffer to receive any access rights sent along with the message is specified in `msg_accrights`, which has length `msg_accrightslen`.

#### RETURN VALUES

These calls return the number of bytes received, or `-1` if an error occurred.

#### ERRORS

EBADF	<code>s</code> is an invalid descriptor.
EFAULT	The data were specified to be received into a non-existent or protected part of the process address space.
EINTR	The calling process received a signal before any data were available to be received, and the signal was set to interrupt the system call.
ENOTSOCK	<code>s</code> is a descriptor for a file, not a socket.
EWOULDBLOCK	The socket is marked non-blocking and the requested operation would block.

#### SEE ALSO

`connect(2)`, `fcntl(2V)`, `getsockopt(2)`, `ioctl(2)`, `read(2V)`, `select(2)`, `send(2)`, `socket(2)`

**NAME**

rename – change the name of a file

**SYNOPSIS**

```
int rename(path1, path2)
char *path1, *path2;
```

**DESCRIPTION**

**rename()** renames the link named *path1* as *path2*. If *path2* exists, then it is first removed. If *path2* refers to a directory, it must be an empty directory, and must not include *path1* in its path prefix. Both *path1* and *path2* must be of the same type (that is, both directories or both non-directories), and must reside on the same file system. Write access permission is required for both the directory containing *path1* and the directory containing *path2*. If a rename request relocates a directory in the hierarchy, write permission in the directory to be moved is needed, since its entry for the parent directory (..) must be updated.

**rename()** guarantees that an instance of *path2* will always exist, even if the system should crash in the middle of the operation.

If the final component of *path1* is a symbolic link, the symbolic link is renamed, not the file or directory to which it points.

If the file referred to by *path2* exists and the file's link count becomes zero when it is removed and no process has the file open, the space occupied by the file is freed, and the file is no longer accessible. If one or more processes have the file open when the last link is removed, the link is removed before **rename()** returns, but the file's contents are not removed until all references to the file have been closed.

Upon successful completion, **rename()** marks for update the *st\_ctime* and *st\_mtime* fields of the parent directory of each file.

**RETURN VALUES**

**rename()** returns:

- 0 on success.
- 1 on failure and sets *errno* to indicate the error.

**ERRORS**

**rename()** will fail and neither *path1* nor *path2* will be affected if:

EACCES	Write access is denied for either <i>path1</i> or <i>path2</i> . A component of the path prefix of either <i>path1</i> or <i>path2</i> denies search permission. The requested rename requires writing in a directory with access permissions that deny write permission.
EBUSY	<i>path2</i> is a directory and is the mount point for a mounted file system.
EDQUOT	The directory in which the entry for the new name is being placed cannot be extended because the user's quota of disk blocks on the file system containing the directory has been exhausted.
EFAULT	Either or both of <i>path1</i> or <i>path2</i> point outside the process's allocated address space.
EINVAL	<i>path1</i> is a parent directory of <i>path2</i> . An attempt was made to rename '.' or '..'.
EIO	An I/O error occurred while reading from or writing to the file system.
EISDIR	<i>path2</i> points to a directory and <i>path1</i> points to a file that is not a directory.
ELOOP	Too many symbolic links were encountered while translating either <i>path1</i> or <i>path2</i> .

ENAMETOOLONG	The length of either path argument exceeds {PATH_MAX}. A pathname component is longer than {NAME_MAX} while {_POSIX_NO_TRUNC} is in effect (see <b>pathconf(2V)</b> ).
ENOENT	A component of the path prefix of either <i>path1</i> or <i>path2</i> does not exist. The file named by <i>path1</i> does not exist.
ENOSPC	The directory in which the entry for the new name is being placed cannot be extended because there is no space left on the file system containing the directory.
ENOTDIR	A component of the path prefix of either <i>path1</i> or <i>path2</i> is not a directory. <i>path1</i> names a directory and <i>path2</i> names a nondirectory file.
ENOTEMPTY	<i>path2</i> is a directory and is not empty.
EROFS	The requested rename requires writing in a directory on a read-only file system.
EXDEV	The link named by <i>path2</i> and the file named by <i>path1</i> are on different logical devices (file systems).

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

ENOENT            *path1* or *path2* points to an empty string.

**SEE ALSO**

**open(2V)**

**WARNINGS**

The system can deadlock if a loop in the file system graph is present. This loop takes the form of an entry in directory **a**, say **a/file1**, being a hard link to directory **b**, and an entry in directory **b**, say **b/file2**, being a hard link to directory **a**. When such a loop exists and two separate processes attempt to perform 'rename **a/file1 b/file2**' and 'rename **b/file2 a/file1**', respectively, the system may deadlock attempting to lock both directories for modification. Hard links to directories should not be used. System administrators should use symbolic links instead.

## NAME

`rmdir` – remove a directory file

## SYNOPSIS

```
int rmdir(path)
char *path;
```

## DESCRIPTION

`rmdir()` removes a directory file whose name is given by *path*. The directory must not have any entries other than `.` and `..`. The directory must not be the root directory or the current directory of the calling process.

If the directory's link count becomes zero, and no process has the directory open, the space occupied by the directory is freed and the directory is no longer accessible. If one or more processes have the directory open when the last link is removed, the `.` and `..` entries, if present, are removed before `rmdir()` returns and no new entries may be created in the directory, but the directory is not removed until all references to the directory have been closed.

Upon successful completion, `rmdir()` marks for update the `st_ctime` and `st_mtime` fields of the parent directory.

## RETURN VALUES

`rmdir()` returns:

- 0 on success.
- 1 on failure and sets `errno` to indicate the error.

## ERRORS

EACCES	Search permission is denied for a component of the path prefix of <i>path</i> .
EACCES	Write permission is denied for the parent directory of the directory to be removed.
EBUSY	The directory to be removed is the mount point for a mounted file system, or is being used by another process.
EFAULT	<i>path</i> points outside the process's allocated address space.
EINVAL	The directory referred to by <i>path</i> is the current directory, <code>.</code> .
EIO	An I/O error occurred while reading from or writing to the file system.
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
ENAMETOOLONG	The length of the path argument exceeds <code>{PATH_MAX}</code> . A pathname component is longer than <code>{NAME_MAX}</code> while <code>{_POSIX_NO_TRUNC}</code> is in effect (see <code>pathconf(2V)</code> ).
ENOENT	The directory referred to by <i>path</i> does not exist.
ENOTDIR	A component of the path prefix of <i>path</i> is not a directory.
ENOTDIR	The file referred to by <i>path</i> is not a directory.
ENOTEMPTY	The directory referred to by <i>path</i> contains files other than <code>.</code> and <code>..</code> .
EROFS	The directory to be removed resides on a read-only file system.

## SYSTEM V ERRORS

In addition to the above, the following may also occur:

- ENOENT *path* points to a null pathname.

## SEE ALSO

`mkdir(2V)`, `unlink(2V)`

**NAME**

select – synchronous I/O multiplexing

**SYNOPSIS**

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

int select (width, readfds, writefds, exceptfds, timeout)
int width;
fd_set *readfds, *writefds, *exceptfds;
struct timeval *timeout;

FD_SET (fd, &fdset)
FD_CLR (fd, &fdset)
FD_ISSET (fd, &fdset)
FD_ZERO (&fdset)
int fd;
fd_set fdset;
```

**DESCRIPTION**

**select()** examines the I/O descriptor sets whose addresses are passed in *readfds*, *writefds*, and *exceptfds* to see if some of their descriptors are ready for reading, ready for writing, or have an exceptional condition pending. *width* is the number of bits to be checked in each bit mask that represent a file descriptor; the descriptors from 0 through *width*-1 in the descriptor sets are examined. Typically *width* has the value returned by **ulimit(3C)** for the maximum number of file descriptors. On return, **select()** replaces the given descriptor sets with subsets consisting of those descriptors that are ready for the requested operation. The total number of ready descriptors in all the sets is returned.

The descriptor sets are stored as bit fields in arrays of integers. The following macros are provided for manipulating such descriptor sets: **FD\_ZERO (&fdset)** initializes a descriptor set *fdset* to the null set. **FD\_SET(*fd*, &fdset)** includes a particular descriptor *fd* in *fdset*. **FD\_CLR(*fd*, &fdset)** removes *fd* from *fdset*. **FD\_ISSET(*fd*, &fdset)** is nonzero if *fd* is a member of *fdset*, zero otherwise. The behavior of these macros is undefined if a descriptor value is less than zero or greater than or equal to **FD\_SETSIZE**, which is normally at least equal to the maximum number of descriptors supported by the system.

If *timeout* is not a NULL pointer, it specifies a maximum interval to wait for the selection to complete. If *timeout* is a NULL pointer, the select blocks indefinitely. To effect a poll, the *timeout* argument should be a non-NULL pointer, pointing to a zero-valued **timeval** structure.

Any of *readfds*, *writefds*, and *exceptfds* may be given as NULL pointers if no descriptors are of interest.

Selecting true for reading on a socket descriptor upon which a **listen(2)** call has been performed indicates that a subsequent **accept(2)** call on that descriptor will not block.

**RETURN VALUES**

**select()** returns a non-negative value on success. A positive value indicates the number of ready descriptors in the descriptor sets. 0 indicates that the time limit referred to by *timeout* expired. On failure, **select()** returns -1, sets **errno** to indicate the error, and the descriptor sets are not changed.

**ERRORS**

EBADF	One of the descriptor sets specified an invalid descriptor.
EFAULT	One of the pointers given in the call referred to a non-existent portion of the process' address space.
EINTR	A signal was delivered before any of the selected events occurred, or before the time limit expired.
EINVAL	A component of the pointed-to time limit is outside the acceptable range: <b>t_sec</b> must be between 0 and $10^8$ , inclusive. <b>t_usec</b> must be greater than or equal to 0, and less than $10^6$ .

**SEE ALSO**

**accept(2), connect(2), fcntl(2V), ulimit(3C), gettimeofday(2), listen(2), read(2V), recv(2), send(2), write(2V)**

**NOTES**

Under rare circumstances, **select()** may indicate that a descriptor is ready for writing when in fact an attempt to write would block. This can happen if system resources necessary for a write are exhausted or otherwise unavailable. If an application deems it critical that writes to a file descriptor not block, it should set the descriptor for non-blocking I/O using the **F\_SETFL** request to **fcntl(2V)**.

**BUGS**

Although the provision of **ulimit(3C)** was intended to allow user programs to be written independent of the kernel limit on the number of open files, the dimension of a sufficiently large bit field for **select** remains a problem. The default size **FD\_SETSIZE** (currently 256) is somewhat larger than the current kernel limit to the number of open files. However, in order to accommodate programs which might potentially use a larger number of open files with **select**, it is possible to increase this size within a program by providing a larger definition of **FD\_SETSIZE** before the inclusion of **<sys/types.h>**.

**select()** should probably return the time remaining from the original timeout, if any, by modifying the time value in place. This may be implemented in future versions of the system. Thus, it is unwise to assume that the timeout pointer will be unmodified by the **select()** call.

## NAME

semctl – semaphore control operations

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/sem.h>

int semctl(semid, semnum, cmd, arg)
int semid, semnum, cmd;
union semun {
    val;
    struct semid_ds *buf;
    ushort *array;
} arg;
```

## DESCRIPTION

**semctl()** provides a variety of semaphore control operations as specified by *cmd*.

The following *cmds* are executed with respect to the semaphore specified by *semid* and *semnum*:

**GETVAL** Return the value of *semval* (see **intro(2)**). [READ]

**SETVAL** Set the value of *semval* to *arg.val*. [ALTER] When this cmd is successfully executed, the *semadj* value corresponding to the specified semaphore in all processes is cleared.

**GETPID** Return the value of *sempid*. [READ]

**GETNCNT** Return the value of *semncnt*. [READ]

**GETZCNT** Return the value of *semzcnt*. [READ]

The following *cmds* return and set, respectively, every *semval* in the set of semaphores.

**GETALL** Place *semvals* into the array pointed to by *arg.array*. [READ]

**SETALL** Set *semvals* according to the array pointed to by *arg.array*. [ALTER] When this cmd is successfully executed the *semadj* values corresponding to each specified semaphore in all processes are cleared.

The following *cmds* are also available:

**IPC\_STAT** Place the current value of each member of the data structure associated with *semid* into the structure pointed to by *arg.buf*. The contents of this structure are defined in **intro(2)**. [READ]

**IPC\_SET** Set the value of the following members of the data structure associated with *semid* to the corresponding value found in the structure pointed to by *arg.buf*:

```
sem_perm.uid
sem_perm.gid
sem_perm.mode /* only low 9 bits */
```

This *cmd* can only be executed by a process that has an effective user ID equal to either that of super-user, or to the value of **sem\_perm.cuid** or **sem\_perm.uid** in the data structure associated with *semid*.

**IPC\_RMID** Remove the semaphore identifier specified by *semid* from the system and destroy the set of semaphores and data structure associated with it. This cmd can only be executed by a process that has an effective user ID equal to either that of super-user, or to the value of **sem\_perm.cuid** or **sem\_perm.uid** in the data structure associated with *semid*.



In the `semop(2)` and `semctl(2)` system call descriptions, the permission required for an operation is given as "[token]", where "token" is the type of permission needed interpreted as follows:

00400	Read by user
00200	Alter by user
00060	Read, Alter by group
00006	Read, Alter by others

Read and Alter permissions on a `semid` are granted to a process if one or more of the following are true:

The effective user ID of the process is super-user.

The effective user ID of the process matches `sem_perm.[c]uid` in the data structure associated with `semid` and the appropriate bit of the "user" portion (0600) of `sem_perm.mode` is set.

The effective user ID of the process does not match `sem_perm.[c]uid` and the effective group ID of the process matches `sem_perm.[c]gid` and the appropriate bit of the "group" portion (060) of `sem_perm.mode` is set.

The effective user ID of the process does not match `sem_perm.[c]uid` and the effective group ID of the process does not match `sem_perm.[c]gid` and the appropriate bit of the "other" portion (06) of `sem_perm.mode` is set.

Otherwise, the corresponding permissions are denied.

#### RETURN VALUES

On success, the value returned by `semctl()` depends on `cmd` as follows:

GETVAL	The value of <code>semval</code> .
GETPID	The value of <code>sempid</code> .
GETNCNT	The value of <code>semncnt</code> .
GETZCNT	The value of <code>semzcnt</code> .
All others	0.

On failure, `semctl()` returns -1 and sets `errno` to indicate the error.

#### ERRORS

EACCES	Operation permission is denied to the calling process (see <code>intro(2)</code> ).
EFAULT	<code>arg.buf</code> points to an illegal address.
EINVAL	<code>semid</code> is not a valid semaphore identifier.  <code>semnum</code> is less than zero or greater than <code>sem_nsems</code> .  <code>cmd</code> is not a valid command.
EPERM	<code>cmd</code> is <code>IPC_RMID</code> or <code>IPC_SET</code> and the effective user ID of the calling process is not super-user.  <code>cmd</code> is <code>IPC_RMID</code> or <code>IPC_SET</code> and the effective user ID of the calling process is not the value of <code>sem_perm.cuid</code> or <code>sem_perm.uid</code> in the data structure associated with <code>semid</code> .
ERANGE	<code>cmd</code> is <code>SETVAL</code> or <code>SETALL</code> and the value to which <code>semval</code> is to be set is greater than the system imposed maximum.

#### SEE ALSO

`intro(2)`, `semget(2)`, `semop(2)`, `ipcrm(1)`, `ipcs(1)`

## NAME

semget – get set of semaphores

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/sem.h>

int semget(key, nsems, semflg)
key_t key;
int nsems, semflg;
```

## DESCRIPTION

semget() returns the semaphore identifier associated with *key*.

A semaphore identifier and associated data structure and set containing *nsems* semaphores (see [intro\(2\)](#)) are created for *key* if one of the following are true:

- *key* is equal to `IPC_PRIVATE`.
- *key* does not already have a semaphore identifier associated with it, and (*semflg* & `IPC_CREAT`) is “true”.

Upon creation, the data structure associated with the new semaphore identifier is initialized as follows:

- `sem_perm.cuid`, `sem_perm.uid`, `sem_perm.cgid`, and `sem_perm.gid` are set equal to the effective user ID and effective group ID, respectively, of the calling process.
- The low-order 9 bits of `sem_perm.mode` are set equal to the low-order 9 bits of *semflg*.
- `sem_nsems` is set equal to the value of *nsems*.
- `sem_otime` is set equal to 0 and `sem_ctime` is set equal to the current time.

A semaphore identifier (*semid*) is a unique positive integer created by a [semget\(2\)](#) system call. Each *semid* has a set of semaphores and a data structure associated with it. The data structure is referred to as `semid_ds` and contains the following members:

```
struct ipc_perm sem_perm; /* operation permission struct */
ushort sem_nsems;        /* number of sems in set */
time_t  sem_otime;       /* last operation time */
time_t  sem_ctime;       /* last change time */
/* Times measured in secs since */
/* 00:00:00 GMT, Jan. 1, 1970 */
```

`sem_perm` is an `ipc_perm` structure that specifies the semaphore operation permission (see below). This structure includes the following members:

```
ushort cuid;             /* creator user id */
ushort cgid;             /* creator group id */
ushort uid;              /* user id */
ushort gid;              /* group id */
ushort mode;             /* r/a permission */
```

The value of `sem_nsems` is equal to the number of semaphores in the set. Each semaphore in the set is referenced by a positive integer referred to as a `sem_num`. `sem_num` values run sequentially from 0 to the value of `sem_nsems` minus 1. `sem_otime` is the time of the last [semop\(2\)](#) operation, and `sem_ctime` is the time of the last [semctl\(2\)](#) operation that changed a member of the above structure.

A semaphore is a data structure that contains the following members:

```

ushort  semval;           /* semaphore value */
short   sempid;          /* pid of last operation */
ushort  semncnt;         /* # awaiting semval > cval */
ushort  semzcnt;         /* # awaiting semval = 0 */

```

**semval** is a non-negative integer. **sempid** is equal to the process ID of the last process that performed a semaphore operation on this semaphore. **semncnt** is a count of the number of processes that are currently suspended awaiting this semaphore's **semval** to become greater than its current value. **semzcnt** is a count of the number of processes that are currently suspended awaiting this semaphore's **semval** to become zero.

#### RETURN VALUES

**semget()** returns a non-negative semaphore identifier on success. On failure, it returns **-1** and sets **errno** to indicate the error.

#### ERRORS

**EACCES** A semaphore identifier exists for *key*, but operation permission (see **intro(2)**) as specified by the low-order 9 bits of *semflg* would not be granted.

**EEXIST** A semaphore identifier exists for *key* but ( (*semflg* & **IPC\_CREAT**) and (*semflg* & **IPC\_EXCL**)) is "true".

**EINVAL** *nsems* is either less than or equal to zero or greater than the system-imposed limit.  
A semaphore identifier exists for *key*, but the number of semaphores in the set associated with it is less than *nsems* and *nsems* is not equal to zero.

**ENOENT** A semaphore identifier does not exist for *key* and (*semflg* & **IPC\_CREAT**) is "false".

**ENOSPC** A semaphore identifier is to be created but the system-imposed limit on the maximum number of allowed semaphore identifiers system wide would be exceeded.  
A semaphore identifier is to be created but the system-imposed limit on the maximum number of allowed semaphores system wide would be exceeded.

#### SEE ALSO

**ipcrm(1)**, **ipcs(1)**, **intro(2)**, **semctl(2)**, **semop(2)**

## NAME

semop – semaphore operations

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/sem.h>

int semop(semid, sops, nsops)
int semid;
struct sembuf *sops;
int nsops;
```

## DESCRIPTION

**semop()** is used to perform atomically an array of semaphore operations on the set of semaphores associated with the semaphore identifier specified by *semid*. *sops* is a pointer to the array of semaphore-operation structures. *nsops* is the number of such structures in the array. The contents of each structure includes the following members:

```
short  sem_num; /* semaphore number */
short  sem_op;  /* semaphore operation */
short  sem_flg; /* operation flags */
```

Each semaphore operation specified by **sem\_op** is performed on the corresponding semaphore specified by *semid* and *sem\_num*.

**sem\_op** specifies one of three semaphore operations as follows:

If **sem\_op** is a negative integer, one of the following will occur: [ALTER] (see **semctl(2)**)

- If *semval* (see **intro(2)**) is greater than or equal to the absolute value of **sem\_op()**, the absolute value of **sem\_op()** is subtracted from *semval*. Also, if (**sem\_flg & SEM\_UNDO**) is “true”, the absolute value of **sem\_op()** is added to the calling process’s *semadj* value (see **exit(2V)**) for the specified semaphore.
- If *semval* is less than the absolute value of **sem\_op()** and (**sem\_flg & IPC\_NOWAIT**) is “true”, **semop()** will return immediately.
- If *semval* is less than the absolute value of **sem\_op()** and (**sem\_flg & IPC\_NOWAIT**) is “false”, **semop()** will increment the *semncnt* associated with the specified semaphore and suspend execution of the calling process until one of the following conditions occur.

*semval* becomes greater than or equal to the absolute value of **sem\_op()**. When this occurs, the value of *semncnt* associated with the specified semaphore is decremented, the absolute value of **sem\_op()** is subtracted from *semval* and, if (**sem\_flg & SEM\_UNDO**) is “true”, the absolute value of **sem\_op()** is added to the calling process’s *semadj* value for the specified semaphore.

The *semid* for which the calling process is awaiting action is removed from the system (see **semctl(2)**). When this occurs, **errno** is set equal to EIDRM, and a value of -1 is returned.

The calling process receives a signal that is to be caught. When this occurs, the value of *semncnt* associated with the specified semaphore is decremented, and the calling process resumes execution in the manner prescribed in **signal(3V)**.

If **sem\_op()** is a positive integer, the value of **sem\_op()** is added to *semval* and, if (**sem\_flg & SEM\_UNDO**) is “true”, the value of **sem\_op()** is subtracted from the calling process’s *semadj* value for the specified semaphore. [ALTER]

If `sem_op()` is zero, one of the following will occur: [READ]

- If `semval` is zero, `semop()` will return immediately.
- If `semval` is not equal to zero and `(sem_flg & IPC_NOWAIT)` is “true”, `semop()` will return immediately.
- If `semval` is not equal to zero and `(sem_flg & IPC_NOWAIT)` is “false”, `semop()` will increment the `semzcnt` associated with the specified semaphore and suspend execution of the calling process until one of the following occurs:
  - `semval` becomes zero, at which time the value of `semzcnt` associated with the specified semaphore is decremented.
  - The `semid` for which the calling process is awaiting action is removed from the system. When this occurs, `errno` is set equal to `EIDRM`, and a value of `-1` is returned.
  - The calling process receives a signal that is to be caught. When this occurs, the value of `semzcnt` associated with the specified semaphore is decremented, and the calling process resumes execution in the manner prescribed in `signal(3V)`.

Upon successful completion, the value of `sempid` for each semaphore specified in the array pointed to by `sops` is set equal to the process ID of the calling process.

#### RETURN VALUES

`semop()` returns:

- 0 on success.
- 1 on failure and sets `errno` to indicate the error.

#### ERRORS

E2BIG	<code>nsops</code> is greater than the system-imposed maximum.
EACCES	Operation permission is denied to the calling process (see <code>intro(2)</code> ).
EAGAIN	The operation would result in suspension of the calling process but <code>(sem_flg &amp; IPC_NOWAIT)</code> is “true”.
EFAULT	<code>sops</code> points to an illegal address.
EFBIG	<code>sem_num</code> is less than zero or greater than or equal to the number of semaphores in the set associated with <code>semid</code> .
EIDRM	The set of semaphores referred to by <code>msqid</code> was removed from the system.
EINTR	The call was interrupted by the delivery of a signal.
EINVAL	<code>semid</code> is not a valid semaphore identifier.
	The number of individual semaphores for which the calling process requests a <code>SEM_UNDO</code> would exceed the limit.
ENOSPC	The limit on the number of individual processes requesting an <code>SEM_UNDO</code> would be exceeded.
ERANGE	An operation would cause a <code>semval</code> or <code>semudj</code> value to overflow the system-imposed limit.

#### SEE ALSO

`ipcrm(1)`, `ipcs(1)`, `intro(2)`, `execve(2V)`, `exit(2V)`, `fork(2V)`, `semctl(2)`, `semget(2)`, `signal(3V)`

## NAME

send, sendto, sendmsg – send a message from a socket

## SYNOPSIS

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

int send(s, msg, len, flags)
int s;
char *msg;
int len, flags;

int sendto(s, msg, len, flags, to, tolen)
int s;
char *msg;
int len, flags;
struct sockaddr *to;
int tolen;

int sendmsg(s, msg, flags)
int s;
struct msghdr *msg;
int flags;
```

## DESCRIPTION

*s* is a socket created with `socket(2)`. `send()`, `sendto()`, and `sendmsg()` are used to transmit a message to another socket. `send()` may be used only when the socket is in a *connected* state, while `sendto()` and `sendmsg()` may be used at any time.

The address of the target is given by *to* with *tolen* specifying its size. The length of the message is given by *len*. If the message is too long to pass atomically through the underlying protocol, then the error EMSGSIZE is returned, and the message is not transmitted.

No indication of failure to deliver is implicit in a `send()`. Return values of `-1` indicate some locally detected errors.

If no buffer space is available at the socket to hold the message to be transmitted, then `send()` normally blocks, unless the socket has been placed in non-blocking I/O mode. The `select(2)` call may be used to determine when it is possible to send more data.

If the process calling `send()`, `sendmsg()` or `sendto()` receives a signal before any data are buffered to be sent, the system call is restarted unless the calling process explicitly set the signal to interrupt these calls using `sigvec()` or `sigaction()` (see the discussions of `SV_INTERRUPT` on `sigvec(2)`, and `SA_INTERRUPT` on `sigaction(3V)`).

The *flags* parameter is formed by ORing one or more of the following:

**MSG\_OOB** Send “out-of-band” data on sockets that support this notion. The underlying protocol must also support “out-of-band” data. Currently, only `SOCK_STREAM` sockets created in the `AF_INET` address family support out-of-band data.

**MSG\_DONTROUTE** The `SO_DONTROUTE` option is turned on for the duration of the operation. This is usually used only by diagnostic or routing programs.

See `recv(2)` for a description of the `msghdr` structure.

## RETURN VALUES

On success, these functions return the number of bytes sent. On failure, they return `-1` and set `errno` to indicate the error.

**ERRORS**

EBADF	<i>s</i> is an invalid descriptor.
EFAULT	The data was specified to be sent to a non-existent or protected part of the process address space.
EINTR	The calling process received a signal before any data could be buffered to be sent, and the signal was set to interrupt the system call.
EINVAL	<i>len</i> is not the size of a valid address for the specified address family.
EMSGSIZE	The socket requires that message be sent atomically, and the size of the message to be sent made this impossible.
ENOBUFS	The system was unable to allocate an internal buffer. The operation may succeed when buffers become available.
ENOBUFS	The output queue for a network interface was full. This generally indicates that the interface has stopped sending, but may be caused by transient congestion.
ENOTSOCK	<i>s</i> is a descriptor for a file, not a socket.
EWOULDBLOCK	The socket is marked non-blocking and the requested operation would block.

**SEE ALSO**

**connect(2), fcntl(2V), getsockopt(2), recv(2), select(2), socket(2), write(2V)**

**NAME**

setpgid – set process group ID for job control

**SYNOPSIS**

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

```
int setpgid (pid, pgid)
```

```
pid_t pid, pgid;
```

**DESCRIPTION**

`setpgid()` is used to either join an existing process group or create a new process group within the session of the calling process (see **NOTES**). The process group ID of a session leader does not change. Upon successful completion, the process group ID of the process with a process ID that matches *pid* is set to *pgid*. As a special case, if *pid* is zero, the process ID of the calling process is used. Also, if *pgid* is zero, the process ID of the process indicated by *pid* is used.

**RETURN VALUES**

`setpgid()` returns:

0       on success.

-1       on failure and sets `errno` to indicate the error.

**ERRORS**

EACCES	The value of <i>pid</i> matches the process ID of a child process of the calling process and the child process has successfully executed one of the <code>exec()</code> functions.
EINVAL	The value of <i>pgid</i> is less than zero or is greater than <code>MAXPID</code> , the maximum process ID as defined in <code>&lt;sys/param.h&gt;</code> .
EPERM	The process indicated by <i>pid</i> is a session leader. The value of <i>pid</i> is valid but matches the process ID of a child process of the calling process and the child process is not in the same session as the calling process. The value of <i>pgid</i> does not match the process ID of the process indicated by <i>pid</i> and there is no process with a process group ID that matches the value of <i>pgid</i> in the same session as the calling process.
ESRCH	<i>pid</i> does not match the PID of the calling process or the PID of a child of the calling process.

**SEE ALSO**

`getpgrp(2V)`, `execve(2V)`, `setsid(2V)`, `tcgetpgrp(3V)`

**NOTES**

For `setpgid()` to behave as described above, `{_POSIX_JOB_CONTROL}` must be in effect (see `sysconf(2V)`). `{_POSIX_JOB_CONTROL}` is always in effect on SunOS systems, but for portability, applications should call `sysconf()` to determine whether `{_POSIX_JOB_CONTROL}` is in effect for the current system.



**NAME**

**setregid** – set real and effective group IDs

**SYNOPSIS**

```
int setregid(rgid, egid)  
int rgid, egid;
```

**DESCRIPTION**

**setregid()** is used to set the real and effective group IDs of the calling process. If *rgid* is  $-1$ , the real GID is not changed; if *egid* is  $-1$ , the effective GID is not changed. The real and effective GIDs may be set to different values in the same call.

If the effective user ID of the calling process is super-user, the real GID and the effective GID can be set to any legal value.

If the effective user ID of the calling process is not super-user, either the real GID can be set to the saved setGID from **execve(2V)**, or the effective GID can either be set to the saved setGID or the real GID. Note: if a setGID process sets its effective GID to its real GID, it can still set its effective GID back to the saved setGID.

In either case, if the real GID is being changed (that is, if *rgid* is not  $-1$ ), or the effective GID is being changed to a value not equal to the real GID, the saved setGID is set equal to the new effective GID.

**RETURN VALUES**

**setregid()** returns:

- 0        on success.
- $-1$      on failure and sets **errno** to indicate the error.

**ERRORS**

**setregid()** will fail and neither of the group IDs will be changed if:

- EINVAL**        The value of *rgid* or *egid* is less than 0 or greater than **USHRT\_MAX** (defined in **<sys/limits.h>**).
- EPERM**        The calling process' effective UID is not the super-user and a change other than changing the real GID to the saved setGID, or changing the effective GID to the real GID or the saved GID, was specified.

**SEE ALSO**

**execve(2V)**, **getgid(2V)**, **setreuid(2)**, **setuid(3V)**

**NAME**

setreuid – set real and effective user IDs

**SYNOPSIS**

```
int setreuid(ruid, euid)
int ruid, euid;
```

**DESCRIPTION**

**setreuid()** is used to set the real and effective user IDs of the calling process. If *ruid* is  $-1$ , the real user ID is not changed; if *euid* is  $-1$ , the effective user ID is not changed. The real and effective user IDs may be set to different values in the same call.

If the effective user ID of the calling process is super-user, the real user ID and the effective user ID can be set to any legal value.

If the effective user ID of the calling process is not super-user, either the real user ID can be set to the effective user ID, or the effective user ID can either be set to the saved set-user ID from **execve(2V)** or the real user ID. Note: if a set-UID process sets its effective user ID to its real user ID, it can still set its effective user ID back to the saved set-user ID.

In either case, if the real user ID is being changed (that is, if *ruid* is not  $-1$ ), or the effective user ID is being changed to a value not equal to the real user ID, the saved set-user ID is set equal to the new effective user ID.

**RETURN VALUES**

**setreuid()** returns:

0        on success.

$-1$       on failure and sets **errno** to indicate the error.

**ERRORS**

**setreuid()** will fail and neither of the user IDs will be changed if:

**EINVAL**        The value of *ruid* or *euid* is less than 0 or greater than **USHRT\_MAX** (defined in **<sys/limits.h>**).

**EPERM**         The calling process' effective user ID is not the super-user and a change other than changing the real user ID to the effective user ID, or changing the effective user ID to the real user ID or the saved set-user ID, was specified.

**SEE ALSO**

**execve(2V)**, **getuid(2V)**, **setregid(2)**, **setuid(3V)**

**NAME**

setsid – create session and set process group ID

**SYNOPSIS**

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

```
pid_t setsid()
```

**DESCRIPTION**

If the calling process is not a process group leader, the `setsid()` function creates a new session. The calling process is the session leader of this new session, the process group leader of a new process group, and has no controlling terminal. If the process had a controlling terminal, `setsid()` breaks the association between the process and that controlling terminal. The process group ID of the calling process is set equal to the process ID of the calling process. The calling process is the only process in the new process group and the only process in the new session.

**RETURN VALUES**

`setsid()` returns the process group ID of the calling process on success. On failure, it returns `-1` and sets `errno` to indicate the error.

**ERRORS**

If any of the following conditions occur, `setsid()` returns `-1` and sets `errno` to the corresponding value:

**EPERM**           The calling process is already a process group leader.

                  The process ID of the calling process equals the process group ID of a different process.

**SEE ALSO**

`execve(2V)`, `exit(2V)`, `fork(2V)`, `getpid(2V)`, `getpgrp(2V)`, `kill(2V)`, `setpgid(2V)`, `sigaction(3V)`

**NAME**

setuseraudit, setaudit – set the audit classes for a specified user ID

**SYNOPSIS**

```
#include <sys/label.h>
#include <sys/audit.h>

int setuseraudit(uid, state)
int uid;
audit_state_t *state;

int setaudit(state)
audit_state_t *state;
```

**DESCRIPTION**

The `setuseraudit()` system call sets the audit state for all processes whose audit user ID matches the specified user ID. The parameter `state` specifies the audit classes to audit for both successful and unsuccessful operations.

The `setaudit()` system call sets the audit state for the current process.

Only processes with the real or effective user ID of the super-user may successfully execute these calls.

**RETURN VALUES**

`setuseraudit()` and `setaudit()` return:

- 0       on success.
- 1       on failure and set `errno` to indicate the error.

**ERRORS**

- EFAULT       The `state` parameter points outside the processes' allocated address space.
- EPERM        The process' real or effective user ID is not super-user.

**SEE ALSO**

`audit(2)`, `audit_args(3)`, `audit_control(5)`, `audit.log(5)`

## NAME

shmctl – shared memory control operations

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/shm.h>

int shmctl (shmid, cmd, buf)
int shmid, cmd;
struct shmctl *buf;
```

## DESCRIPTION

shmctl() provides a variety of shared memory control operations as specified by *cmd*. The following *cmds* are available:

- IPC\_STAT** Place the current value of each member of the data structure associated with *shmid* into the structure pointed to by *buf*. The contents of this structure are defined in **intro(2)**. [READ]
- IPC\_SET** Set the value of the following members of the data structure associated with *shmid* to the corresponding value found in the structure pointed to by *buf*:

```
shm_perm.uid
shm_perm.gid
shm_perm.mode /* only low 9 bits */
```

This *cmd* can only be executed by a process that has an effective user ID equal to that of super-user, or to the value of **shm\_perm.cuid** or **shm\_perm.uid** in the data structure associated with *shmid*.

- IPC\_RMID** Remove the shared memory identifier specified by *shmid* from the system. If no processes are currently mapped to the corresponding shared memory segment, then the segment is removed and the associated resources are reclaimed. Otherwise, the segment will persist, although **shmget(2)** will not be able to locate it, until it is no longer mapped by any process. This *cmd* can only be executed by a process that has an effective user ID equal to that of super-user, or to the value of **shm\_perm.cuid** or **shm\_perm.uid** in the data structure associated with *shmid*.

In the **shmop(2)** and **shmctl(2)** system call descriptions, the permission required for an operation is given as "[token]", where "token" is the type of permission needed interpreted as follows:

00400	Read by user
00200	Write by user
00060	Read, Write by group
00006	Read, Write by others

Read and Write permissions on a *shmid* are granted to a process if one or more of the following are true:

The effective user ID of the process is super-user.

The effective user ID of the process matches **shm\_perm.[c]uid** in the data structure associated with *shmid* and the appropriate bit of the "user" portion (0600) of **shm\_perm.mode** is set.

The effective user ID of the process does not match **shm\_perm.[c]uid** and the effective group ID of the process matches **shm\_perm.[c]gid** and the appropriate bit of the "group" portion (060) of **shm\_perm.mode** is set.

The effective user ID of the process does not match `shm_perm.[c]uid` and the effective group ID of the process does not match `shm_perm.[c]gid` and the appropriate bit of the “other” portion (06) of `shm_perm.mode` is set.

Otherwise, the corresponding permissions are denied.

#### RETURN VALUES

`shmctl()` returns:

- 0 on success.
- 1 on failure and sets `errno` to indicate the error.

#### ERRORS

- EACCES** `cmd` is equal to `IPC_STAT` and `[READ]` operation permission is denied to the calling process (see `intro(2)`).
- EFAULT** `buf` points to an illegal address.
- EINVAL** `shmid` is not a valid shared memory identifier.  
`cmd` is not a valid command.
- EPERM** `cmd` is equal to `IPC_RMID` or `IPC_SET` and the effective user ID of the calling process is not super-user or the value of `shm_perm.cuid` or `shm_perm.uid` in the data structure associated with `shmid`.

#### SEE ALSO

`ipcrm(1)`, `ipcs(1)`, `intro(2)`, `shmget(2)`, `shmop(2)`

## NAME

shmget – get shared memory segment identifier

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/shm.h>

int shmget(key, size, shmflg)
key_t key;
int size, shmflg;
```

## DESCRIPTION

shmget() returns the shared memory identifier associated with *key*.

A shared memory identifier and associated data structure and shared memory segment of at least *size* bytes (see intro(2)) are created for *key* if one of the following are true:

- *key* is equal to IPC\_PRIVATE.
- *key* does not already have a shared memory identifier associated with it, and (*shmflg* & IPC\_CREAT) is “true”.

Upon creation, the data structure associated with the new shared memory identifier is initialized as follows:

- *shm\_perm.cuid*, *shm\_perm.uid*, *shm\_perm.cgid*, and *shm\_perm.gid* are set equal to the effective user ID and effective group ID, respectively, of the calling process.
- The low-order 9 bits of *shm\_perm.mode* are set equal to the low-order 9 bits of *shmflg*.
- *shm\_segsz* is set equal to the value of *size*.
- *shm\_lpid*, *shm\_nattch*, *shm\_atime*, and *shm\_dtime* are set equal to 0.
- *shm\_ctime* is set equal to the current time.

A shared memory identifier (shmid) is a unique positive integer created by a shmget(2) system call. Each shmid has a segment of memory (referred to as a shared memory segment) and a data structure associated with it. The data structure is referred to as *shmid\_ds* and contains the following members:

```
struct ipc_perm shm_perm; /* operation permission struct */
int shm_segsz; /* size of segment */
ushort shm_cpid; /* creator pid */
ushort shm_lpid; /* pid of last operation */
short shm_nattch; /* number of current attaches */
time_t shm_atime; /* last attach time */
time_t shm_dtime; /* last detach time */
time_t shm_ctime; /* last change time */
/* Times measured in secs since */
/* 00:00:00 GMT, Jan. 1, 1970 */
```

*shm\_perm* is an *ipc\_perm* structure that specifies the shared memory operation permission (see below). This structure includes the following members:

```
ushort cuid; /* creator user id */
ushort cgid; /* creator group id */
ushort uid; /* user id */
ushort gid; /* group id */
ushort mode; /* r/w permission */
```

**shm\_segsz** specifies the size of the shared memory segment. **shm\_cpid** is the process ID of the process that created the shared memory identifier. **shm\_lpid** is the process ID of the last process that performed a **shmop(2)** operation. **shm\_nattch** is the number of processes that currently have this segment attached. **shm\_atime** is the time of the last **shmat** operation, **shm\_dtime** is the time of the last **shmdt** operation, and **shm\_ctime** is the time of the last **shmctl(2)** operation that changed one of the members of the above structure.

#### RETURN VALUES

**shmget()** returns a non-negative shared memory identifier on success. On failure, it returns **-1** and sets **errno** to indicate the error.

#### ERRORS

EACCES	A shared memory identifier exists for <i>key</i> but operation permission (see <b>intro(2)</b> ) as specified by the low-order 9 bits of <i>shmflg</i> would not be granted.
EEXIST	A shared memory identifier exists for <i>key</i> but ( ( <i>shmflg</i> & <b>IPC_CREAT</b> ) && ( <i>shmflg</i> & <b>IPC_EXCL</b> ) ) is "true".
EINVAL	<i>size</i> is less than the system-imposed minimum or greater than the system-imposed maximum. A shared memory identifier exists for <i>key</i> but the size of the segment associated with it is less than <i>size</i> and <i>size</i> is not equal to zero.
ENOENT	A shared memory identifier does not exist for <i>key</i> and ( <i>shmflg</i> & <b>IPC_CREAT</b> ) is "false".
ENOMEM	A shared memory identifier and associated shared memory segment are to be created but the amount of available physical memory is not sufficient to fill the request.
ENOSPC	A shared memory identifier is to be created but the system-imposed limit on the maximum number of allowed shared memory identifiers system wide would be exceeded.

#### SEE ALSO

**ipcrm(1)**, **ipcs(1)**, **intro(2)**, **shmctl(2)**, **shmop(2)**



**NAME**

shmop, shmat, shmdt – shared memory operations

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/shm.h>

char *shmat(shmid, shmaddr, shmflg)
int shmid;
char *shmaddr;
int shmflg;

int shmdt(shmaddr)
char *shmaddr;
```

**DESCRIPTION**

**shmat()** maps the shared memory segment associated with the shared memory identifier specified by *shmid* into the data segment of the calling process. Upon successful completion, the address of the mapped segment is returned.

The shared memory segment is mapped at the address specified by one of the following criteria:

- If *shmaddr* is equal to zero, the segment is mapped at an address selected by the system. Ordinarily, applications should invoke **shmat()** with *shmaddr* equal to zero so that the operating system may make the best use of available resources.
- If *shmaddr* is not equal to zero and (*shmflg* & SHM\_RND) is “true”, the segment is mapped at the address given by (*shmaddr* - (*shmaddr* modulus SHMLBA)).
- If *shmaddr* is not equal to zero and (*shmflg* & SHM\_RND) is “false”, the segment is mapped at the address given by *shmaddr*.

The segment is mapped for reading if (*shmflg* & SHM\_RDONLY) is “true” [READ], otherwise it is mapped for reading and writing [READ/WRITE] (see **shmctl(2)**).

**shmdt()** unmaps from the calling process’s address space the shared memory segment that is mapped at the address specified by *shmaddr*. The shared memory segment must have been mapped with a prior **shmat()** function call. The segment and contents are retained until explicitly removed by means of the IPC\_RMID function (see **shmctl(2)**).

**RETURN VALUES**

**shmat()** returns the data segment start address of the mapped shared memory segment. On failure, it returns -1 and sets **errno** to indicate the error.

**shmdt()** returns:

- 0 on success.
- 1 on failure and sets **errno** to indicate the error.

**ERRORS**

**shmat()** will fail and not map the shared memory segment if one or more of the following are true:

- |        |  |
|--------|--|
| EACCES | Operation permission is denied to the calling process (see <b>intro(2)</b> ).  |
| EINVAL | <i>shmid</i> is not a valid shared memory identifier.<br><br><i>shmaddr</i> is not equal to zero, and the value of ( <i>shmaddr</i> - ( <i>shmaddr</i> modulus SHMLBA)) is an illegal address.<br><br><i>shmaddr</i> is not equal to zero, ( <i>shmflg</i> & SHM_RND) is “false”, and the value of <i>shmaddr</i> is an illegal address. |
| EMFILE | The number of shared memory segments mapped to the calling process would exceed the system-imposed limit.  |

**ENOMEM**        The available data space is not large enough to accommodate the shared memory segment.

**shmdt()** will fail and not unmap the shared memory segment if:

**EINVAL**        *shmaddr* is not the data segment start address of a shared memory segment.

**SEE ALSO**

**ipcrm(1), ipcs(1), intro(2), execve(2V), exit(2V), fork(2V), shmctl(2), shmget(2)**

**NAME**

shutdown – shut down part of a full-duplex connection

**SYNOPSIS**

```
int shutdown(s, how)
int s, how;
```

**DESCRIPTION**

The **shutdown()** call causes all or part of a full-duplex connection on the socket associated with *s* to be shut down. If *how* is 0, then further receives will be disallowed. If *how* is 1, then further sends will be disallowed. If *how* is 2, then further sends and receives will be disallowed.

**RETURN VALUES**

**shutdown()** returns:

0 on success.

-1 on failure and sets **errno** to indicate the error.

**ERRORS**

**EBADF** *s* is not a valid descriptor.

**ENOTCONN** The specified socket is not connected.

**ENOTSOCK** *s* is a file, not a socket.

**SEE ALSO**

**ipcrm(1)**, **ipcs(1)**, **connect(2)**, **socket(2)**

**BUGS**

The *how* values should be defined constants.

**NAME**

sigblock, sigmask – block signals

**SYNOPSIS**

```
#include <signal.h>
```

```
int sigblock(mask);
```

```
int mask;
```

```
int sigmask(signum)
```

**DESCRIPTION**

**sigblock()** adds the signals specified in *mask* to the set of signals currently being blocked from delivery. A signal is blocked if the appropriate bit in *mask* is set. The macro **sigmask()** is provided to construct the signal mask for a given *signum*. **sigblock()** returns the previous signal mask, which may be restored using **sigsetmask(2)**.

It is not possible to block SIGKILL or SIGSTOP. The system silently imposes this restriction.

**RETURN VALUES**

**sigblock()** returns the previous signal mask.

The **sigmask()** macro returns the mask for the given signal number.

**SEE ALSO**

**kill(2V)**, **sigsetmask(2)**, **sigvec(2)**, **signal(3V)**

**NAME**

**sigpause**, **sigsuspend** – automatically release blocked signals and wait for interrupt

**SYNOPSIS**

```
int sigpause(sigmask)
int sigmask;

#include <signal.h>

int sigsuspend(sigmaskp)
sigset_t *sigmaskp;
```

**DESCRIPTION**

**sigpause()** assigns *sigmask* to the set of masked signals and then waits for a signal to arrive; on return the set of masked signals is restored. *sigmask* is usually 0 to indicate that no signals are now to be blocked. **sigpause()** always terminates by being interrupted, returning EINTR.

In normal usage, a signal is blocked using **sigblock(2)**, to begin a critical section, variables modified on the occurrence of the signal are examined to determine that there is no work to be done, and the process pauses awaiting work by using **sigpause()** with the mask returned by **sigblock()**.

**sigsuspend()** replaces the process's signal mask with the set of signals pointed to by *sigmaskp* and then suspends the process until delivery of a signal whose action is either to execute a signal-catching function or to terminate the process. If the action is to terminate the process, **sigsuspend()** does not return. If the action is to execute a signal-catching function, **sigsuspend()** returns after the signal-catching function returns, with the signal mask restored to the setting that existed prior to the **sigsuspend()** call. It is not possible to block those signals that cannot be ignored, as documented in <signal.h> this is enforced by the system without indicating an error.

**RETURN VALUES**

Since **sigpause()** and **sigsuspend()** suspend process execution indefinitely, there is no successful completion return value. On failure, these functions return -1 and set **errno** to indicate the error.

**ERRORS**

EINTR            A signal is caught by the calling process and control is returned from the signal-catching function.

**SEE ALSO**

**sigblock(2)**, **sigpending(2V)**, **sigprocmask(2V)**, **sigvec(2)**, **pause(3V)**, **sigaction(3V)**, **signal(3V)**, **sigsetops(3V)**

**NAME**

sigpending – examine pending signals

**SYNOPSIS**

```
#include <signal.h>
```

```
int sigpending(set)
```

```
sigset_t *set;
```

**DESCRIPTION**

**sigpending()** stores the set of signals that are blocked from delivery and pending for the calling process in the space pointed to by *set*.

**RETURN VALUES**

**sigpending()** returns:

0        on success.

-1       on failure and sets **errno** to indicate the error.

**SEE ALSO**

**sigprocmask(2V)**, **sigvec(2)**, **sigsetops(3V)**

**NAME**

`sigprocmask` – examine and change blocked signals

**SYNOPSIS**

```
#include <signal.h>

int sigprocmask(how, set, oset)
int how;
sigset_t *set, *oset;
```

**DESCRIPTION**

`sigprocmask()` is used to examine or change (or both) the calling process's signal mask. If the value of *set* is not NULL, it points to a set of signals to be used to change the currently blocked set.

The value of *how* indicates the manner in which the set is changed, and consists of one of the following values, as defined in the header `<signal.h>`:

- SIG\_BLOCK**     The resulting set is the union of the current set and the signal set pointed to by *set*.
- SIG\_UNBLOCK**   The resulting set is the intersection of the current set and the complement of the signal set pointed to by *set*.
- SIG\_SETMASK**   The resulting set is the signal set pointed to by *set*.

If *oset* is not NULL, the previous mask is stored in the space pointed to by *oset*. If the value of *set* is NULL, the value of *how* is not significant and the process's signal mask is unchanged by this function call. Thus, the call can be used to enquire about currently blocked signals.

If there are any pending unblocked signals after the call to `sigprocmask()`, at least one of those signals is delivered before `sigprocmask()` returns.

If it is not possible to block the **SIGKILL** and **SIGSTOP** signals. This is enforced by the system without causing an error to be indicated.

If any of the **SIGFPE**, **SIGKILL**, or **SIGSEGV** signals are generated while they are blocked, the result is undefined, unless the signal was generated by a call to `kill(2V)`.

If `sigprocmask()` fails, the process's signal mask is not changed.

**RETURN VALUES**

`sigprocmask()` returns:

- 0           on success.
- 1          on failure and sets `errno` to indicate the error.

**ERRORS**

- EINVAL**           The value of *how* is not equal to one of the defined values.

**SEE ALSO**

`sigpause(2V)`, `sigpending(2V)`, `sigvec(2)`, `sigaction(3V)`, `sigsetops(3V)`

**NAME**

sigsetmask – set current signal mask

**SYNOPSIS**

```
#include <signal.h>
int sigsetmask(mask)
int mask;
```

**DESCRIPTION**

sigsetmask() sets the set of signals currently being blocked from delivery according to *mask*. A signal is blocked if the appropriate bit in *mask* is set. The macro sigblock(2) is provided to construct the mask for a given *signum*.

The system silently disallows blocking SIGKILL and SIGSTOP.

**RETURN VALUES**

sigsetmask() returns the previous signal mask.

**SEE ALSO**

kill(2V), sigblock(2), sigpause(2V), sigvec(2), signal(3V)



**NAME**

sigstack – set and/or get signal stack context

**SYNOPSIS**

```
#include <signal.h>
```

```
int sigstack (ss, oss)
```

```
struct sigstack *ss, *oss;
```

**DESCRIPTION**

**sigstack()** allows users to define an alternate stack, called the “signal stack”, on which signals are to be processed. When a signal’s action indicates its handler should execute on the signal stack (specified with a **sigvec(2)** call), the system checks to see if the process is currently executing on that stack. If the process is not currently executing on the signal stack, the system arranges a switch to the signal stack for the duration of the signal handler’s execution.

A signal stack is specified by a **sigstack()** structure, which includes the following members:

```
char          *ss_sp;          /* signal stack pointer */
int           ss_onstack;     /* current status */
```

**ss\_sp** is the initial value to be assigned to the stack pointer when the system switches the process to the signal stack. Note that, on machines where the stack grows downwards in memory, this is *not* the address of the beginning of the signal stack area. **ss\_onstack** field is zero or non-zero depending on whether the process is currently executing on the signal stack or not.

If **ss** is not a NULL pointer, **sigstack()** sets the signal stack state to the value in the **sigstack()** structure pointed to by **ss**. Note: if **ss\_onstack** is non-zero, the system will think that the process is executing on the signal stack. If **ss** is a NULL pointer, the signal stack state will be unchanged. If **oss** is not a NULL pointer, the current signal stack state is stored in the **sigstack()** structure pointed to by **oss**.

**RETURN VALUES**

**sigstack()** returns:

0           on success.

-1          on failure and sets **errno** to indicate the error.

**ERRORS**

**sigstack()** will fail and the signal stack context will remain unchanged if one of the following occurs.

**EFAULT**           **ss** or **oss** points to memory that is not a valid part of the process address space.

**SEE ALSO**

**sigvec(2)**, **setjmp(3V)**, **signal(3V)**

**NOTES**

Signal stacks are not “grown” automatically, as is done for the normal stack. If the stack overflows unpredictable results may occur.

## NAME

sigvec – software signal facilities

## SYNOPSIS

```
#include <signal.h>

int sigvec(sig, vec, ovec)
int sig;
struct sigvec *vec, *ovec;
```

## DESCRIPTION

The system defines a set of signals that may be delivered to a process. Signal delivery resembles the occurrence of a hardware interrupt: the signal is blocked from further occurrence, the current process context is saved, and a new one is built. A process may specify a *handler* to which a signal is delivered, or specify that a signal is to be *blocked* or *ignored*. A process may also specify that a default action is to be taken by the system when a signal occurs. Normally, signal handlers execute on the current stack of the process. This may be changed, on a per-handler basis, so that signals are taken on a special *signal stack*.

All signals have the same *priority*. Signal routines execute with the signal that caused their invocation *blocked*, but other signals may yet occur. A global *signal mask* defines the set of signals currently blocked from delivery to a process. The signal mask for a process is initialized from that of its parent (normally 0). It may be changed with a `sigblock(2)` or `sigsetmask(2)` call, or when a signal is delivered to the process.

A process may also specify a set of *flags* for a signal that affect the delivery of that signal.

When a signal condition arises for a process, the signal is added to a set of signals pending for the process. If the signal is not currently *blocked* by the process then it is delivered to the process. When a signal is delivered, the current state of the process is saved, a new signal mask is calculated (as described below), and the signal handler is invoked. The call to the handler is arranged so that if the signal handling routine returns normally the process will resume execution in the context from before the signal's delivery. If the process wishes to resume in a different context, then it must arrange to restore the previous context itself.

When a signal is delivered to a process a new signal mask is installed for the duration of the process' signal handler (or until a `sigblock()` or `sigsetmask()` call is made). This mask is formed by taking the current signal mask, adding the signal to be delivered, and ORing in the signal mask associated with the handler to be invoked.

The action to be taken when the signal is delivered is specified by a `sigvec` structure, defined in `<signal.h>` as:

```
struct sigvec {
    void (*sv_handler)();    /* signal handler */
    int sv_mask;            /* signal mask to apply */
    int sv_flags;           /* see signal options */
}
```

The following bits may be set in `sv_flags`:

```
#define SV_ONSTACK      0x0001    /* take signal on signal stack */
#define SV_INTERRUPT    0x0002    /* do not restart system on signal return */
#define SV_RESETHAND    0x0004    /* reset signal handler to SIG_DFL on signal */
```

If the `SV_ONSTACK` bit is set in the flags for that signal, the system will deliver the signal to the process on the signal stack specified with `sigstack(2)`, rather than delivering the signal on the current stack.

If `vec` is not a NULL pointer, `sigvec()` assigns the handler specified by `sv_handler`, the mask specified by `sv_mask`, and the flags specified by `sv_flags` to the specified signal. If `vec` is a NULL pointer, `sigvec()` does not change the handler, mask, or flags for the specified signal.

The mask specified in `vec` is not allowed to block `SIGKILL` or `SIGSTOP`. The system enforces this restriction silently.

If *ovec* is not a NULL pointer, the handler, mask, and flags in effect for the signal before the call to `sigvec()` are returned to the user. A call to `sigvec()` with *vec* a NULL pointer and *ovec* not a NULL pointer can be used to determine the handling information currently in effect for a signal without changing that information.

The following is a list of all signals with names as in the include file `<signal.h>`:

<b>SIGHUP</b>	1	hangup
<b>SIGINT</b>	2	interrupt
<b>SIGQUIT</b>	3*	quit
<b>SIGILL</b>	4*	illegal instruction
<b>SIGTRAP</b>	5*	trace trap
<b>SIGABRT</b>	6*	abort (generated by <code>abort(3)</code> routine)
<b>SIGEMT</b>	7*	emulator trap
<b>SIGFPE</b>	8*	arithmetic exception
<b>SIGKILL</b>	9	kill (cannot be caught, blocked, or ignored)
<b>SIGBUS</b>	10*	bus error
<b>SIGSEGV</b>	11*	segmentation violation
<b>SIGSYS</b>	12*	bad argument to system call
<b>SIGPIPE</b>	13	write on a pipe or other socket with no one to read it
<b>SIGALRM</b>	14	alarm clock
<b>SIGTERM</b>	15	software termination signal
<b>SIGURG</b>	16●	urgent condition present on socket
<b>SIGSTOP</b>	17†	stop (cannot be caught, blocked, or ignored)
<b>SIGTSTP</b>	18†	stop signal generated from keyboard
<b>SIGCONT</b>	19●	continue after stop
<b>SIGCHLD</b>	20●	child status has changed
<b>SIGTTIN</b>	21†	background read attempted from control terminal
<b>SIGTTOU</b>	22†	background write attempted to control terminal
<b>SIGIO</b>	23●	I/O is possible on a descriptor (see <code>fcntl(2V)</code> )
<b>SIGXCPU</b>	24	cpu time limit exceeded (see <code>getrlimit(2)</code> )
<b>SIGXFSZ</b>	25	file size limit exceeded (see <code>getrlimit(2)</code> )
<b>SIGVTALRM</b>	26	virtual time alarm (see <code>getitimer(2)</code> )
<b>SIGPROF</b>	27	profiling timer alarm (see <code>getitimer(2)</code> )
<b>SIGWINCH</b>	28●	window changed (see <code>termio(4)</code> and <code>win(4S)</code> )
<b>SIGLOST</b>	29*	resource lost (see <code>lockd(8C)</code> )
<b>SIGUSR1</b>	30	user-defined signal 1
<b>SIGUSR2</b>	31	user-defined signal 2

The starred signals in the list above cause a core image if not caught or ignored.

Once a signal handler is installed, it remains installed until another `sigvec()` call is made, or an `execve(2V)` is performed, unless the `SV_RESETHAND` bit is set in the flags for that signal. In that case, the value of the handler for the caught signal is set to `SIG_DFL` before entering the signal-catching function, unless the signal is `SIGILL` or `SIGTRAP`. Also, if this bit is set, the bit for that signal in the signal mask will not be set; unless the signal mask associated with that signal blocks that signal, further occurrences of that signal will not be blocked. The `SV_RESETHAND` flag is not available in 4.2BSD, hence it should not be used if backward compatibility is needed.

The default action for a signal may be reinstated by setting the signal's handler to `SIG_DFL`; this default is termination except for signals marked with ● or †. Signals marked with ● are discarded if the action is `SIG_DFL`; signals marked with † cause the process to stop. If the process is terminated, a "core image" will be made in the current working directory of the receiving process if the signal is one for which an asterisk appears in the above list *and* the following conditions are met:

- The effective user ID (EUID) and the real user ID (UID) of the receiving process are equal.
- The effective group ID (EGID) and the real group ID (GID) of the receiving process are equal.

- An ordinary file named `core` exists and is writable or can be created. If the file must be created, it will have the following properties:
  - a mode of 0666 modified by the file creation mask (see `umask(2V)`)
  - a file owner ID that is the same as the effective user ID of the receiving process.
  - a file group ID that is the same as the file group ID of the current directory

If the handler for that signal is `SIG_IGN`, the signal is subsequently ignored, and pending instances of the signal are discarded.

Note: the signals `SIGKILL` and `SIGSTOP` cannot be ignored.

If a caught signal occurs during certain system calls, the call is restarted by default. The call can be forced to terminate prematurely with an `EINTR` error return by setting the `SV_INTERRUPT` bit in the flags for that signal. `SV_INTERRUPT` is not available in 4.2BSD, hence it should not be used if backward compatibility is needed. The affected system calls are `read(2V)` or `write(2V)` on a slow device (such as a terminal or pipe or other socket, but not a file) and during a `wait(2V)`.

After a `fork(2V)`, or `vfork(2)` the child inherits all signals, the signal mask, the signal stack, and the restart/interrupt and reset-signal-handler flags.

The `execve(2V)`, call resets all caught signals to default action and resets all signals to be caught on the user stack. Ignored signals remain ignored; the signal mask remains the same; signals that interrupt system calls continue to do so.

## CODES

The following defines the codes for signals which produce them. All of these symbols are defined in `signal.h`:

Condition	Signal	Code
<b>Sun codes:</b>		
Illegal instruction	<code>SIGILL</code>	<code>ILL_INSTR_FAULT</code>
Integer division by zero	<code>SIGFPE</code>	<code>FPE_INTDIV_TRAP</code>
IEEE floating pt inexact	<code>SIGFPE</code>	<code>FPE_FLTINEX_TRAP</code>
IEEE floating pt division by zero	<code>SIGFPE</code>	<code>FPE_FLTDIV_TRAP</code>
IEEE floating pt underflow	<code>SIGFPE</code>	<code>FPE_FLTUND_TRAP</code>
IEEE floating pt operand error	<code>SIGFPE</code>	<code>FPE_FLTOPERR_TRAP</code>
IEEE floating pt overflow	<code>SIGFPE</code>	<code>FPE_FLTOVF_FAULT</code>
Hardware bus error	<code>SIGBUS</code>	<code>BUS_HWERR</code>
Address alignment error	<code>SIGBUS</code>	<code>BUS_ALIGN</code>
No mapping fault	<code>SIGSEGV</code>	<code>SEGV_NOMAP</code>
Protection fault	<code>SIGSEGV</code>	<code>SEGV_PROT</code>
Object error	<code>SIGSEGV</code>	<code>SEGV_CODE(code)=SEGV_OBJERR</code>
Object error number	<code>SIGSEGV</code>	<code>SEGV_ERRNO(code)</code>
<b>SPARC codes:</b>		
Privileged instruction violation	<code>SIGILL</code>	<code>ILL_PRIVINSTR_FAULT</code>
Bad stack	<code>SIGILL</code>	<code>ILL_STACK</code>
Trap # <i>n</i> (1 <= <i>n</i> <= 127)	<code>SIGILL</code>	<code>ILL_TRAP_FAULT(<i>n</i>)</code>
Integer overflow	<code>SIGFPE</code>	<code>FPE_INTOVF_TRAP</code>
Tag overflow	<code>SIGEMT</code>	<code>EMT_TAG</code>
<b>MC680X0 codes:</b>		
Privilege violation	<code>SIGILL</code>	<code>ILL_PRIVVIO_FAULT</code>
Coprocessor protocol error	<code>SIGILL</code>	<code>ILL_INSTR_FAULT</code>
Trap # <i>n</i> (1 <= <i>n</i> <= 14)	<code>SIGILL</code>	<code>ILL_TRAP<i>n</i>_FAULT</code>
A-line op code	<code>SIGEMT</code>	<code>EMT_EMU1010</code>
F-line op code	<code>SIGEMT</code>	<code>EMT_EMU1111</code>
CHK or CHK2 instruction	<code>SIGFPE</code>	<code>FPE_CHKINST_TRAP</code>
TRAPV or TRAPcc or cpTRAPcc	<code>SIGFPE</code>	<code>FPE_TRAPV_TRAP</code>

IEEE floating pt compare unordered	<b>SIGFPE</b>	<b>FPE_FLTBSUN_TRAP</b>
IEEE floating pt signaling NaN	<b>SIGFPE</b>	<b>FPE_FLTNAN_TRAP</b>

**ADDR**

The *addr* signal handler parameter is defined as follows:

Signal	Code	Addr
Sun:		
<b>SIGILL</b>	Any	address of faulted instruction
<b>SIGEMT</b>	Any	address of faulted instruction
<b>SIGFPE</b>	Any	address of faulted instruction
<b>SIGBUS</b>	<b>BUS_HWERR</b>	address that caused fault
<b>SIGSEGV</b>	Any	address that caused fault
SPARC:		
<b>SIGBUS</b>	<b>BUS_ALIGN</b>	address of faulted instruction
MC680X0:		
<b>SIGBUS</b>	<b>BUS_ALIGN</b>	address that caused fault

The accuracy of *addr* is machine dependent. For example, certain machines may supply an address that is on the same page as the address that caused the fault. If an appropriate *addr* cannot be computed it will be set to **SIG\_NOADDR**.

**RETURN VALUES**

**sigvec()** returns:

0        on success.  
 -1       on failure and sets **errno** to indicate the error.

**ERRORS**

**sigvec()** will fail and no new signal handler will be installed if one of the following occurs:

**EFAULT**        Either *vec* or *ovec* is not a NULL pointer and points to memory that is not a valid part of the process address space.  
**EINVAL**        *Sig* is not a valid signal number.

An attempt was made to ignore or supply a handler for **SIGKILL** or **SIGSTOP**.

**SEE ALSO**

**execve(2V)**, **fcntl(2V)**, **fork(2V)**, **getitimer(2)**, **getrlimit(2)**, **ioctl(2)**, **kill(2V)**, **ptrace(2)**, **read(2V)**, **sigblock(2)**, **sigpause(2V)**, **sigsetmask(2)**, **sigstack(2)**, **umask(2V)**, **vfork(2)**, **wait(2V)**, **write(2V)**, **setjmp(3V)**, **signal(3V)**, **streamio(4)**, **termio(4)**, **win(4S)**, **lockd(8C)**

**NOTES**

**SIGPOLL** is a synonym for **SIGIO**. A **SIGIO** will be issued when a file descriptor corresponding to a **STREAMS** (see **intro(2)**) file has a "selectable" event pending. Unless that descriptor has been put into asynchronous mode (see **fcntl(2V)**), a process must specifically request that this signal be sent using the **I\_SETSIG ioctl(2)** call (see **streamio(4)**). Otherwise, the process will never receive **SIGPOLL**.

The handler routine can be declared:

```
void handler(sig, code, scp, addr)
int sig, code;
struct sigcontext *scp;
char *addr;
```

Here *sig* is the signal number; *code* is a parameter of certain signals that provides additional detail; *scp* is a pointer to the `sigcontext` structure (defined in `signal.h`), used to restore the context from before the signal; and *addr* is additional address information.

Programs that must be portable to UNIX systems other than 4.2BSD should use the `signal(3V)`, interface instead.

**NAME**

**socket** – create an endpoint for communication

**SYNOPSIS**

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

int socket(domain, type, protocol)
int domain, type, protocol;
```

**DESCRIPTION**

**socket()** creates an endpoint for communication and returns a descriptor.

The *domain* parameter specifies a communications domain within which communication will take place; this selects the protocol family which should be used. The protocol family generally is the same as the address family for the addresses supplied in later operations on the socket. These families are defined in the include file `<sys/socket.h>`. The currently understood formats are

<b>PF_UNIX</b>	(UNIX system internal protocols),
<b>PF_INET</b>	(ARPA Internet protocols), and
<b>PF_IMPLINK</b>	(IMP “host at IMP” link layer).

The socket has the indicated *type*, which specifies the semantics of communication. Currently defined types are:

```
SOCK_STREAM
SOCK_DGRAM
SOCK_RAW
SOCK_SEQPACKET
SOCK_RDM
```

A **SOCK\_STREAM** type provides sequenced, reliable, two-way connection based byte streams. An out-of-band data transmission mechanism may be supported. A **SOCK\_DGRAM** socket supports datagrams (connectionless, unreliable messages of a fixed (typically small) maximum length). A **SOCK\_SEQPACKET** socket may provide a sequenced, reliable, two-way connection-based data transmission path for datagrams of fixed maximum length; a consumer may be required to read an entire packet with each read system call. This facility is protocol specific, and presently not implemented for any protocol family. **SOCK\_RAW** sockets provide access to internal network interfaces. The types **SOCK\_RAW**, which is available only to the super-user, and **SOCK\_RDM**, for which no implementation currently exists, are not described here.

The *protocol* specifies a particular protocol to be used with the socket. Normally only a single protocol exists to support a particular socket type within a given protocol family. However, it is possible that many protocols may exist, in which case a particular protocol must be specified in this manner. The protocol number to use is particular to the “communication domain” in which communication is to take place; see **protocols(5)**.

Sockets of type **SOCK\_STREAM** are full-duplex byte streams, similar to pipes. A stream socket must be in a *connected* state before any data may be sent or received on it. A connection to another socket is created with a **connect(2)** call. Once connected, data may be transferred using **read(2V)** and **write(2V)** calls or some variant of the **send(2)** and **recv(2)** calls. When a session has been completed a **close(2V)**, may be performed. Out-of-band data may also be transmitted as described in **send(2)** and received as described in **recv(2)**.

The communications protocols used to implement a `SOCK_STREAM` insure that data is not lost or duplicated. If a piece of data for which the peer protocol has buffer space cannot be successfully transmitted within a reasonable length of time, then the connection is considered broken and calls will indicate an error with `-1` returns and with `ETIMEDOUT` as the specific code in the global variable `errno`. The protocols optionally keep sockets “warm” by forcing transmissions roughly every minute in the absence of other activity. An error is then indicated if no response can be elicited on an otherwise idle connection for a extended period (for instance 5 minutes). A `SIGPIPE` signal is raised if a process sends on a broken stream; this causes naive processes, which do not handle the signal, to exit.

`SOCK_SEQPACKET` sockets employ the same system calls as `SOCK_STREAM` sockets. The only difference is that `read(2V)` calls will return only the amount of data requested, and any remaining in the arriving packet will be discarded.

`SOCK_DGRAM` and `SOCK_RAW` sockets allow sending of datagrams to correspondents named in `send(2)` calls. Datagrams are generally received with `recv(2)`, which returns the next datagram with its return address.

An `fcntl(2V)` call can be used to specify a process group to receive a `SIGURG` signal when the out-of-band data arrives. It may also enable non-blocking I/O and asynchronous notification of I/O events with `SIGIO` signals.

The operation of sockets is controlled by socket level *options*. These options are defined in the file `socket.h`. `getsockopt(2)` and `setsockopt()` are used to get and set options, respectively.

#### RETURN VALUES

`socket()` returns a non-negative descriptor on success. On failure, it returns `-1` and sets `errno` to indicate the error.

#### ERRORS

<code>EACCES</code>	Permission to create a socket of the specified type and/or protocol is denied.
<code>EMFILE</code>	The per-process descriptor table is full.
<code>ENFILE</code>	The system file table is full.
<code>ENOBUFS</code>	Insufficient buffer space is available. The socket cannot be created until sufficient resources are freed.
<code>EPROTONOSUPPORT</code>	The protocol type or the specified protocol is not supported within this domain.
<code>EPROTOTYPE</code>	The protocol is the wrong type for the socket.

#### SEE ALSO

`accept(2)`, `bind(2)`, `close(2V)`, `connect(2)`, `fcntl(2V)`, `getsockname(2)`, `getsockopt(2)`, `ioctl(2)`, `listen(2)`, `read(2V)`, `recv(2)`, `select(2)`, `send(2)`, `shutdown(2)`, `socketpair(2)`, `write(2V)`, `protocols(5)`

*Network Programming*



**NAME**

**socketpair** – create a pair of connected sockets

**SYNOPSIS**

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

int socketpair(d, type, protocol, sv)
int d, type, protocol;
int sv[2];
```

**DESCRIPTION**

The **socketpair()** system call creates an unnamed pair of connected sockets in the specified address family *d*, of the specified *type* and using the optionally specified *protocol*. The descriptors used in referencing the new sockets are returned in *sv*[0] and *sv*[1]. The two sockets are indistinguishable.

**RETURN VALUES**

**socketpair()** returns:

- 0        on success.
- 1       on failure and sets **errno** to indicate the error.

**ERRORS**

- EAFNOSUPPORT**        The specified address family is not supported on this machine.
- EFAULT**                The address *sv* does not specify a valid part of the process address space.
- EMFILE**                Too many descriptors are in use by this process.
- EOPNOSUPPORT**        The specified protocol does not support creation of socket pairs.
- EPROTONOSUPPORT**    The specified protocol is not supported on this machine.

**SEE ALSO**

**pipe(2V)**, **read(2V)**, **write(2V)**

**BUGS**

This call is currently implemented only for the **AF\_UNIX** address family.

## NAME

stat, lstat, fstat – get file status

## SYNOPSIS

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

```
#include <sys/stat.h>
```

```
int stat(path, buf)
```

```
char *path;
```

```
struct stat *buf;
```

```
int lstat(path, buf)
```

```
char *path;
```

```
struct stat *buf;
```

```
int fstat(fd, buf)
```

```
int fd;
```

```
struct stat *buf;
```

## DESCRIPTION

**stat()** obtains information about the file named by *path*. Read, write or execute permission of the named file is not required, but all directories listed in the path name leading to the file must be searchable.

**lstat()** is like **stat()** except in the case where the named file is a symbolic link, in which case **lstat()** returns information about the link, while **stat()** returns information about the file the link references.

**fstat()** obtains the same information about an open file referenced by the argument descriptor, such as would be obtained by an **open(2V)** call.

*buf* is a pointer to a **stat** structure into which information is placed concerning the file. A **stat** structure includes the following members:

```

dev_t      st_dev;    /* device file resides on */
ino_t      st_ino;    /* the file serial number */
mode_t     st_mode;   /* file mode */
nlink_t    st_nlink;  /* number of hard links to the file */
uid_t      st_uid;    /* user ID of owner */
gid_t      st_gid;    /* group ID of owner */
dev_t      st_rdev;   /* the device identifier (special files only)*/
off_t      st_size;   /* total size of file, in bytes */
time_t     st_atime;  /* file last access time */
time_t     st_mtime;  /* file last modify time */
time_t     st_ctime;  /* file last status change time */
long       st_blksize; /* preferred blocksize for file system I/O*/
long       st_blocks; /* actual number of blocks allocated */

```

**st\_atime** Time when file data was last accessed. This can also be set explicitly by **utimes(2)**. **st\_atime** is not updated for directories searched during pathname resolution.

**st\_mtime** Time when file data was last modified. This can also be set explicitly by **utimes(2)**. It is not set by changes of owner, group, link count, or mode.

**st\_ctime** Time when file status was last changed. It is set both both by writing and changing the file status information, such as changes of owner, group, link count, or mode.

The following macros test whether a file is of the specified type. The value *m* is the value of **st\_mode**. Each macro evaluates to a non-zero value if the test is true or to zero if the test is false.

**S\_ISDIR(*m*)** Test for directory file.

**S\_ISCHR(*m*)** Test for character special file.

**S\_ISBLK(*m*)** Test for block special file.

<b>S_ISREG(<i>m</i>)</b>	Test for regular file.
<b>S_ISLNK(<i>m</i>)</b>	Test for a symbolic link.
<b>S_ISSOCK(<i>m</i>)</b>	Test for a socket.
<b>S_ISFIFO(<i>m</i>)</b>	Test for pipe or FIFO special file.
The status information word <b>st_mode</b> is bit-encoded using the following masks and bits:	
<b>S_IRWXU</b>	Read, write, search (if a directory), or execute (otherwise) permissions mask for the owner of the file.
<b>S_IRUSR</b>	Read permission bit for the owner of the file.
<b>S_IWUSR</b>	Write permission bit for the owner of the file.
<b>S_IXUSR</b>	Search (if a directory) or execute (otherwise) permission bit for the owner of the file.
<b>S_IRWXG</b>	Read, write, search (if directory), or execute (otherwise) permissions mask for the file group class.
<b>S_IRGRP</b>	Read permission bit for the file group class.
<b>S_IWGRP</b>	Write permission bit for the file group class.
<b>S_IXGRP</b>	Search (if a directory) or execute (otherwise) permission bit for the file group class.
<b>S_IRWXO</b>	Read, write, search (if a directory), or execute (otherwise) permissions mask for the file other class.
<b>S_IROTH</b>	Read permission bit for the file other class.
<b>S_IWOTH</b>	Write permission bit for the file other class.
<b>S_IXOTH</b>	Search (if a directory) or execute (otherwise) permission bit for the file other class.
<b>S_ISUID</b>	Set user ID on execution. The process's effective user ID is set to that of the owner of the file when the file is run as a program (see <code>execve(2V)</code> ). On a regular file, this bit should be cleared on any write.
<b>S_ISGID</b>	Set group ID on execution. The process's effective group ID is set to that of the file when the file is run as a program (see <code>execve(2V)</code> ). On a regular file, this bit should be cleared on any write.

In addition, the following bits and masks are made available for backward compatibility:

<b>#define S_IFMT</b>	<b>0170000</b>	<b>/* type of file */</b>
<b>#define S_IFIFO</b>	<b>0010000</b>	<b>/* FIFO special */</b>
<b>#define S_IFCHR</b>	<b>0020000</b>	<b>/* character special */</b>
<b>#define S_IFDIR</b>	<b>0040000</b>	<b>/* directory */</b>
<b>#define S_IFBLK</b>	<b>0060000</b>	<b>/* block special */</b>
<b>#define S_IFREG</b>	<b>0100000</b>	<b>/* regular file */</b>
<b>#define S_IFLNK</b>	<b>0120000</b>	<b>/* symbolic link */</b>
<b>#define S_IFSOCK</b>	<b>0140000</b>	<b>/* socket */</b>
<b>#define S_ISVTX</b>	<b>0001000</b>	<b>/* save swapped text even after use */</b>
<b>#define S_IREAD</b>	<b>0000400</b>	<b>/* read permission, owner */</b>
<b>#define S_IWRITE</b>	<b>0000200</b>	<b>/* write permission, owner */</b>
<b>#define S_IEXEC</b>	<b>0000100</b>	<b>/* execute/search permission, owner */</b>

For more information on **st\_mode** bits see `chmod(2V)`.

**RETURN VALUES**

**stat()**, **lstat()** and **fstat()** return:

- 0 on success.
- 1 on failure and set **errno** to indicate the error.

**ERRORS**

**stat()** and **lstat()** will fail if one or more of the following are true:

- EACCES** Search permission is denied for a component of the path prefix of *path*.
- EFAULT** *buf* or *path* points to an invalid address.
- EIO** An I/O error occurred while reading from or writing to the file system.
- ELOOP** Too many symbolic links were encountered in translating *path*.
- ENAMETOOLONG** The length of the path argument exceeds **{PATH\_MAX}.n**  
A pathname component is longer than **{NAME\_MAX}** while **{\_POSIX\_NO\_TRUNC}** is in effect (see **pathconf(2V)**).
- ENOENT** The file referred to by *path* does not exist.
- ENOTDIR** A component of the path prefix of *path* is not a directory.

**fstat()** will fail if one or more of the following are true:

- EBADF** *fd* is not a valid open file descriptor.
- EFAULT** *buf* points to an invalid address.
- EIO** An I/O error occurred while reading from or writing to the file system.

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

- ENOENT** *path* points to an empty string.

**WARNINGS**

The **st\_atime** and **st\_mtime** fields of the **stat()** are *not* contiguous. Programs that depend on them being contiguous (in calls to **utimes(2)** or **utime(3V)**) will not work.

**SEE ALSO**

**chmod(2V)**, **chown(2V)**, **link(2V)**, **open(2V)**, **read(2V)**, **readlink(2)**, **rename(2V)**, **truncate(2)**, **unlink(2V)**, **utimes(2)**, **write(2V)**

## NAME

statfs, fstatfs – get file system statistics

## SYNOPSIS

```
#include <sys/vfs.h>

int statfs(path, buf)
char *path;
struct statfs *buf;

int fstatfs(fd, buf)
int fd;
struct statfs *buf;
```

## DESCRIPTION

**statfs()** returns information about a mounted file system. *path* is the path name of any file within the mounted filesystem. *buf* is a pointer to a **statfs()** structure defined as follows:

```
typedef struct {
    long    val[2];
} fsid_t;

struct statfs {
    long    f_type;      /* type of info, zero for now */
    long    f_bsize;     /* fundamental file system block size */
    long    f_blocks;   /* total blocks in file system */
    long    f_bfree;    /* free blocks */
    long    f_bavail;   /* free blocks available to non-super-user */
    long    f_files;    /* total file nodes in file system */
    long    f_ffree;    /* free file nodes in fs */
    fsid_t  f_fsid;     /* file system id */
    long    f_spare[7]; /* spare for later */
};
```

Fields that are undefined for a particular file system are set to -1. **fstatfs()** returns the same information about an open file referenced by descriptor *fd*.

## RETURN VALUES

**statfs()** and **fstatfs()** return:

- 0        on success.
- 1       on failure and set **errno** to indicate the error.

## ERRORS

**statfs()** fails if one or more of the following are true:

- |              |  |
|--------------|--|
| EACCES       | Search permission is denied for a component of the path prefix of <i>path</i> .  |
| EFAULT       | <i>buf</i> or <i>path</i> points to an invalid address.  |
| EIO          | An I/O error occurred while reading from or writing to the file system.  |
| ELOOP        | Too many symbolic links were encountered in translating <i>path</i> .  |
| ENAMETOOLONG | The length of the path argument exceeds {PATH_MAX}.<br>A pathname component is longer than {NAME_MAX} (see <b>sysconf(2V)</b> ) while {_POSIX_NO_TRUNC} is in effect (see <b>pathconf(2V)</b> ). |
| ENOENT       | The file referred to by <i>path</i> does not exist.  |
| ENOTDIR      | A component of the path prefix of <i>path</i> is not a directory.  |

**fstatfs()** fails if one or more of the following are true:

- EBADF                *fd* is not a valid open file descriptor.
- EFAULT              *buf* points to an invalid address.
- EIO                  An I/O error occurred while reading from the file system.

#### BUGS

The NFS revision 2 protocol does not permit the number of free files to be provided to the client; thus, when **statfs()** or **fstatfs()** are done on a file on an NFS file system, **f\_files** and **f\_ffree** are always -1.

**NAME**

swapon – add a swap device for interleaved paging/swapping

**SYNOPSIS**

```
int swapon(special)
char *special;
```

**DESCRIPTION**

**swapon()** makes the block device *special* available to the system for allocation for paging and swapping. The names of potentially available devices are known to the system and defined at system configuration time. The size of the swap area on *special* is calculated at the time the device is first made available for swapping.

**RETURN VALUES**

**swapon()** returns:

- 0        on success.
- 1       on failure and sets **errno** to indicate the error.

**ERRORS**

EACCES	Search permission is denied for a component of the path prefix of <i>special</i> .
EBUSY	The device referred to by <i>special</i> has already been made available for swapping.
EFAULT	<i>special</i> points outside the process's address space.
EIO	An I/O error occurred while reading from or writing to the file system. An I/O error occurred while opening the swap device.
ELOOP	Too many symbolic links were encountered in translating <i>special</i> .
ENAMETOOLONG	The length of the path argument exceeds {PATH_MAX}. A pathname component is longer than {NAME_MAX} (see <b>sysconf(2V)</b> ) while {_POSIX_NO_TRUNC} is in effect (see <b>pathconf(2V)</b> ).
ENODEV	The device referred to by <i>special</i> was not configured into the system as a swap device.
ENOENT	The device referred to by <i>special</i> does not exist.
ENOTBLK	The file referred to by <i>special</i> is not a block device.
ENOTDIR	A component of the path prefix of <i>special</i> is not a directory.
ENXIO	The major device number of the device referred to by <i>special</i> is out of range (this indicates no device driver exists for the associated hardware).
EPERM	The caller is not the super-user.

**SEE ALSO**

**fstab(5)**, **config(8)**, **swapon(8)**

**BUGS**

There is no way to stop swapping on a disk so that the pack may be dismounted.  
This call will be upgraded in future versions of the system.

**NAME**

symlink – make symbolic link to a file

**SYNOPSIS**

```
int symlink(name1, name2)
char *name1, *name2;
```

**DESCRIPTION**

A symbolic link *name2* is created to *name1* (*name2* is the name of the file created, *name1* is the string used in creating the symbolic link). Either name may be an arbitrary path name; the files need not be on the same file system.

The file that the symbolic link points to is used when an `open(2V)` operation is performed on the link. A `stat(2V)`, on a symbolic link returns the linked-to file, while an `lstat()` (refer to `stat(2V)`) returns information about the link itself. This can lead to surprising results when a symbolic link is made to a directory. To avoid confusion in programs, the `readlink(2)` call can be used to read the contents of a symbolic link.

**RETURN VALUES**

`symlink()` returns:

- 0       on success.
- 1       on failure and sets `errno` to indicate the error.

**ERRORS**

The symbolic link is made unless one or more of the following are true:

- |                     |   |
|---------------------|---|
| <b>EACCES</b>       | Search permission is denied for a component of the path prefix of <i>name2</i> .  |
| <b>EDQUOT</b>       | The directory in which the entry for the new symbolic link is being placed cannot be extended because the user's quota of disk blocks on the file system containing the directory has been exhausted.<br><br>The new symbolic link cannot be created because the user's quota of disk blocks on the file system which will contain the link has been exhausted.<br><br>The user's quota of inodes on the file system on which the file is being created has been exhausted. |
| <b>EEXIST</b>       | The file referred to by <i>name2</i> already exists.  |
| <b>EFAULT</b>       | <i>name1</i> or <i>name2</i> points outside the process's allocated address space.  |
| <b>EIO</b>          | An I/O error occurred while reading from or writing to the file system.   |
| <b>ELOOP</b>        | Too many symbolic links were encountered in translating <i>name2</i> .  |
| <b>ENAMETOOLONG</b> | The length of the path argument exceeds <code>{PATH_MAX}</code> .<br><br>A pathname component is longer than <code>{NAME_MAX}</code> (see <code>sysconf(2V)</code> ) while <code>{_POSIX_NO_TRUNC}</code> is in effect (see <code>pathconf(2V)</code> ).  |
| <b>ENOENT</b>       | A component of the path prefix of <i>name2</i> does not exist.  |
| <b>ENOSPC</b>       | The directory in which the entry for the new symbolic link is being placed cannot be extended because there is no space left on the file system containing the directory.<br><br>The new symbolic link cannot be created because there is no space left on the file system which will contain the link.   |
| <b>ENOTDIR</b>      | There are no free inodes on the file system on which the file is being created.<br><br>A component of the path prefix of <i>name2</i> is not a directory.   |
| <b>EROFS</b>        | The file <i>name2</i> would reside on a read-only file system.  |



**SEE ALSO**

**ln(1V), link(2V), readlink(2), unlink(2V)**

**NAME**

sync – update super-block

**SYNOPSIS**

sync()

**DESCRIPTION**

sync() writes out all information in core memory that should be on disk. This includes modified super blocks, modified inodes, and delayed block I/O.

sync() should be used by programs that examine a file system, for example fsck(8), df(1V), etc. sync() is mandatory before a boot.

**SEE ALSO**

fsync(2), cron(8)

**BUGS**

The writing, although scheduled, is not necessarily complete upon return from sync().

**NAME**

`syscall` – indirect system call

**SYNOPSIS**

```
#include <sys/syscall.h>
```

```
int syscall(number[ , arg, ... ] )
```

```
int number;
```

**DESCRIPTION**

`syscall()` performs the system call whose assembly language interface has the specified *number*, and arguments *arg* .... Symbolic constants for system calls can be found in the header file `<sys/syscall.h>`.

**RETURN VALUES**

`syscall()` returns the return value of the system call specified by *number*.

**SEE ALSO**

`intro(2)`, `pipe(2V)`

**WARNINGS**

There is no way to use `syscall()` to call functions such as `pipe(2V)`, which return values that do not fit into one hardware register.

Since many system calls are implemented as library wrappers around traps to the kernel, these calls may not behave as documented when called from `syscall()`, which bypasses these wrappers. For these reasons, using `syscall()` is not recommended.

**NAME**

`sysconf` – query system related limits, values, options

**SYNOPSIS**

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

```
long sysconf(name)
```

```
int name;
```

**DESCRIPTION**

The `sysconf()` function provides a method for the application to determine the current value of a configurable system limit or option (variable). The value does not change during the lifetime of the calling process.

The convention used throughout sections 2 and 3 is that {LIMIT} means that LIMIT is something that can change from system to system and applications that want accurate values need to call `sysconf()`. These values are things that have been historically available in header files such as `<sys/param.h>`.

The following lists the conceptual name and meaning of each variable.

<i>Name</i>	<i>Meaning</i>
{ARG_MAX}	Max combined size of <code>argv[ ]</code> & <code>envp[ ]</code> .
{CHILD_MAX}	Max processes allowed to any UID.
{CLK_TCK}	Ticks per second ( <code>clock_t</code> ).
{NGROUPS_MAX}	Max simultaneous groups one may belong to.
{OPEN_MAX}	Max open files per process.
{_POSIX_JOB_CONTROL}	Job control supported (boolean).
{_POSIX_SAVED_IDS}	Saved ids ( <code>seteuid()</code> ) supported (boolean).
{_POSIX_VERSION}	Version of the POSIX.1 standard supported.

The following table lists the conceptual name of each variable and the flag passed to `sysconf()` to retrieve the value of each variable.

<i>Name</i>	<i>Sysconf flag</i>
{ARG_MAX}	<code>_SC_ARG_MAX</code>
{CHILD_MAX}	<code>_SC_CHILD_MAX</code>
{CLK_TCK}	<code>_SC_CLK_TCK</code>
{NGROUPS_MAX}	<code>_SC_NGROUPS_MAX</code>
{OPEN_MAX}	<code>_SC_OPEN_MAX</code>
{_POSIX_JOB_CONTROL}	<code>_SC_JOB_CONTROL</code>
{_POSIX_SAVED_IDS}	<code>_SC_SAVED_IDS</code>
{_POSIX_VERSION}	<code>_SC_VERSION</code>

**RETURN VALUES**

`sysconf()` returns the current variable value on success. On failure, it returns `-1` and sets `errno` to indicate the error.

**ERRORS**

`EINVAL` The value of *name* is invalid.

## NAME

`truncate`, `ftruncate` – set a file to a specified length

## SYNOPSIS

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

int truncate(path, length)
char *path;
off_t length;

int ftruncate(fd, length)
int fd;
off_t length;
```

## DESCRIPTION

`truncate()` causes the file referred to by *path* (or for `ftruncate()` the object referred to by *fd*) to have a size equal to *length* bytes. If the file was previously longer than *length*, the extra bytes are removed from the file. If it was shorter, bytes between the old and new lengths are read as zeroes. With `ftruncate()`, the file must be open for writing.

## RETURN VALUES

`truncate()` returns:

- 0       on success.
- 1       on failure and sets `errno` to indicate the error.

## ERRORS

`truncate()` may set `errno` to:

- EACCES       Search permission is denied for a component of the path prefix of *path*.  
Write permission is denied for the file referred to by *path*.
- EFAULT       *path* points outside the process's allocated address space.
- EIO           An I/O error occurred while reading from or writing to the file system.
- EISDIR       The file referred to by *path* is a directory.
- ELOOP        Too many symbolic links were encountered in translating *path*.
- ENAMETOOLONG   The length of the path argument exceeds {PATH\_MAX}.  
A pathname component is longer than {NAME\_MAX} (see `sysconf(2V)`) while {\_POSIX\_NO\_TRUNC} is in effect (see `pathconf(2V)`).
- ENOENT       The file referred to by *path* does not exist.
- ENOTDIR       A component of the path prefix of *path* is not a directory.
- EROFS        The file referred to by *path* resides on a read-only file system.

`ftruncate()` may set `errno` to:

- EINVAL       *fd* is not a valid descriptor of a file open for writing.  
*fd* refers to a socket, not to a file.
- EIO           An I/O error occurred while reading from or writing to the file system.

## SEE ALSO

`open(2V)`

## BUGS

These calls should be generalized to allow ranges of bytes in a file to be discarded.

**NAME**

**umask** – set file creation mode mask

**SYNOPSIS**

```
#include <sys/stat.h>
```

```
int umask(mask)
```

```
int mask;
```

**SYSTEM V SYNOPSIS**

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

```
#include <sys/stat.h>
```

```
mode_t umask(mask)
```

```
mode_t mask;
```

**DESCRIPTION**

**umask()** sets the process's file creation mask to *mask* and returns the previous value of the mask. The low-order 9 bits of *mask* are used whenever a file is created, clearing corresponding bits in the file access permissions. (see **stat(2V)**). This clearing restricts the default access to a file.

The mask is inherited by child processes.

**RETURN VALUES**

**umask()** returns the previous value of the file creation mask.

**SEE ALSO**

**chmod(2V)**, **mknod(2V)**, **open(2V)**

**NAME**

`uname` – get information about current system

**SYNOPSIS**

```
#include <sys/utsname.h>
```

```
int uname (name)
```

```
struct utsname *name;
```

**DESCRIPTION**

`uname()` stores information identifying the current operating system in the structure pointed to by *name*.

`uname()` uses the structure defined in `<sys/utsname.h>`, the members of which are:

```
struct utsname {
    char    sysname[9];
    char    nodename[9];
    char    nodeext[65-9];
    char    release[9];
    char    version[9];
    char    machine[9];
}
```

`uname()` places a null-terminated character string naming the current operating system in the character array *sysname*; this string is “SunOS” on Sun systems. *nodename* is set to the name that the system is known by on a communications network; this is the same value as is returned by `gethostname(2)`. *release* and *version* are set to values that further identify the operating system. *machine* is set to a standard name that identifies the hardware on which the SunOS system is running. This is the same as the value displayed by `arch(1)`.

**RETURN VALUES**

`uname()` returns:

0        on success.

-1       on failure.

**SEE ALSO**

`arch(1)`, `uname(1)`, `gethostname(2)`

**NOTES**

*nodeext* is provided for backwards compatibility with previous SunOS Releases and provides space for node names longer than eight bytes. Applications should not use *nodeext*. To be maximally portable, applications that want to copy the node name to another string should use `strlen(nodename)` rather than the constant 9 or `sizeof(nodename)` as the size of the target string.

System administrators should note that systems with node names longer than eight bytes do not conform to *IEEE Std 1003.1-1988, System V Interface Definition (Issue 2)*, or *X/Open Portability Guide (Issue 2)* requirements.

**NAME**

`unlink` – remove directory entry

**SYNOPSIS**

```
int unlink(path)
char *path;
```

**DESCRIPTION**

`unlink()` removes the directory entry named by the pathname pointed to by *path* and decrements the link count of the file referred to by that entry. If this entry was the last link to the file, and no process has the file open, then all resources associated with the file are reclaimed. If, however, the file was open in any process, the actual resource reclamation is delayed until it is closed, even though the directory entry has disappeared.

If *path* refers to a directory, the effective user-ID of the calling process must be super-user.

Upon successful completion, `unlink()` marks for update the `st_ctime` and `st_mtime` fields of the parent directory. Also, if the file's link count is not zero, the `st_ctime` field of the file is marked for update.

**RETURN VALUES**

`unlink()` returns:

- 0        on success.
- 1       on failure and sets `errno` to indicate the error.

**ERRORS**

EACCES	Search permission is denied for a component of the path prefix of <i>path</i> . Write permission is denied for the directory containing the link to be removed.
EBUSY	The entry to be unlinked is the mount point for a mounted file system.
EFAULT	<i>path</i> points outside the process's allocated address space.
EINVAL	The file referred to by <i>path</i> is the current directory, '.'.
EIO	An I/O error occurred while reading from or writing to the file system.
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
ENAMETOOLONG	The length of the path argument exceeds <code>{PATH_MAX}</code> . A pathname component is longer than <code>{NAME_MAX}</code> while <code>{POSIX_NO_TRUNC}</code> is in effect (see <code>pathconf(2V)</code> ).
ENOENT	The file referred to by <i>path</i> does not exist.
ENOTDIR	A component of the path prefix of <i>path</i> is not a directory.
EPERM	The file referred to by <i>path</i> is a directory and the effective user ID of the process is not the super-user.
EROFS	The file referred to by <i>path</i> resides on a read-only file system.

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

- ENOENT        *path* points to an empty string.

**SEE ALSO**

`close(2V)`, `link(2V)`, `rmdir(2V)`

**NOTES**

Applications should use `rmdir(2V)` to remove directories. Although `root` may use `unlink()` on directories, all users may use `rmdir()`.



**NAME**

umount, umount – remove a file system

**SYNOPSIS**

```
int unmount(name)
char *name;
```

**SYSTEM V SYNOPSIS**

```
int umount(special)
char *special;
```

**DESCRIPTION**

**umount()** announces to the system that the directory *name* is no longer to refer to the root of a mounted file system. The directory *name* reverts to its ordinary interpretation.

Only the super-user may call **umount()**.

**SYSTEM V DESCRIPTION**

**umount()** requests that a previously mounted file system contained on the block special device referred to by *special* be unmounted. *special* points to a path name. After the file system is unmounted, the directory on which it was mounted reverts to its ordinary interpretation.

Only the super-user may call **umount()**.

Note: Unlike the path name argument to **umount()** which refers to the directory on which the file system is mounted, *special* refers to the block special device containing the mounted file system itself.

**RETURN VALUES**

**umount()** returns:

- 0 on success.
- 1 on failure and sets **errno** to indicate the error.

**SYSTEM V RETURN VALUES**

**umount()** returns:

- 0 on success.
- 1 on failure and sets **errno** to indicate the error.

**ERRORS**

EACCES	Search permission is denied for a component of the path prefix.
EBUSY	A process is holding a reference to a file located on the file system.
EFAULT	<i>name</i> points outside the process's allocated address space.
EINVAL	<i>name</i> is not the root of a mounted file system.
EIO	An I/O error occurred while reading from or writing to the file system.
ELOOP	Too many symbolic links were encountered in translating the path name.
ENAMETOOLONG	The length of the path argument exceeds {PATH_MAX}.
	A pathname component is longer than {NAME_MAX} (see <b>sysconf(2V)</b> ) while {_POSIX_NO_TRUNC} is in effect (see <b>pathconf(2V)</b> ).
ENOENT	<i>name</i> does not exist.
ENOTDIR	A component of the path prefix of <i>name</i> is not a directory.
EPERM	The caller is not the super-user.

**SYSTEM V ERRORS**

EINVAL	The device referred to by <i>special</i> is not mounted.
ENOENT	The named file does not exist.

ENOTBLK            *special* does not refer to a block special file.  
ENOTDIR            A component of the path prefix of *special* is not a directory.  
ENXIO              The device referred to by *special* does not exist.

**SEE ALSO**

**mount(2V), mount(8)**

**BUGS**

The error codes are in a state of disarray; too many errors appear to the caller as one value.

## NAME

`ustat` – get file system statistics

## SYNOPSIS

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

```
#include <ustat.h>
```

```
int ustat(dev, buf)
```

```
dev_t dev;
```

```
struct ustat *buf;
```

## DESCRIPTION

`ustat()` returns information about a mounted file system. *dev* is a device number identifying a device containing a mounted file system. This is normally the value returned in the `st_dev` field of a `stat` structure when a `stat()`, `fstat()`, or `lstat()` call is made on a file on that file system. *buf* is a pointer to a `ustat` structure that includes the following elements:

```
    daddr_t f_tfree;           /* Total blocks available to non-super-user */
```

```
    ino_t   f_tinode;        /* Number of free files */
```

```
    char    f_fname[6];     /* Filsys name */
```

```
    char    f_fpack[6];     /* Filsys pack name */
```

The `f_fname` and `f_fpack` fields are always set to a null string. Other fields that are undefined for a particular file system are set to `-1`.

## RETURN VALUES

`ustat()` returns:

0        on success.

-1       on failure and sets `errno` to indicate the error.

## ERRORS

EFAULT        *buf* points to an invalid address.

EINVAL        *dev* is not the device number of a device containing a mounted file system.

EIO            An I/O error occurred while reading from or writing to the file system.

## SEE ALSO

`stat(2V)`, `statfs(2)`

## BUGS

The NFS revision 2 protocol does not permit the number of free files to be provided to the client; thus, when `ustat()` is done on an NFS file system, `f_tinode` is always `-1`.

**NAME**

utimes – set file times

**SYNOPSIS**

```
#include <sys/types.h>
int utimes(file, tvp)
char *file;
struct timeval *tvp;
```

**DESCRIPTION**

**utimes()** sets the access and modification times of the file named by *file*.

If *tvp* is NULL, the access and modification times are set to the current time. A process must be the owner of the file or have write permission for the file to use **utimes()** in this manner.

If *tvp* is not NULL, it is assumed to point to an array of two **timeval** structures. The access time is set to the value of the first member, and the modification time is set to the value of the second member. Only the owner of the file or the super-user may use **utimes()** in this manner.

In either case, the *inode-changed* time of the file is set to the current time.

**RETURN VALUES**

**utimes()** returns:

- 0 on success.
- 1 on failure and sets **errno** to indicate the error.

**ERRORS**

- EACCES** Search permission is denied for a component of the path prefix of *file*.
- EACCES** The effective user ID of the process is not super-user and not the owner of the file, write permission is denied for the file, and *tvp* is NULL.
- EFAULT** *file* or *tvp* points outside the process's allocated address space.
- EIO** An I/O error occurred while reading from or writing to the file system.
- ELOOP** Too many symbolic links were encountered in translating *file*.
- ENOENT** The file referred to by *file* does not exist.
- ENOTDIR** A component of the path prefix of *file* is not a directory.
- EPERM** The effective user ID of the process is not super-user and not the owner of the file, and *tvp* is not NULL.
- EROFS** The file system containing the file is mounted read-only.

**SEE ALSO**

**stat(2V)**

**NAME**

vadvise – give advice to paging system

**SYNOPSIS**

```
#include <sys/vadvise.h>
```

```
vadvise(param)
```

```
int param;
```

**DESCRIPTION**

**vadvise()** is used to inform the system that process paging behavior merits special consideration. Parameters to **vadvise()** are defined in the file `<sys/vadvise.h>`. Currently, two calls to **vadvise()** are implemented.

```
vadvise(VA_ANOM);
```

advises that the paging behavior is not likely to be well handled by the system's default algorithm, since reference information that is collected over macroscopic intervals (for instance, 10-20 seconds) will not serve to indicate future page references. The system in this case will choose to replace pages with little emphasis placed on recent usage, and more emphasis on referenceless circular behavior. It is *essential* that processes which have very random paging behavior (such as LISP during garbage collection of very large address spaces) call **vadvise**, as otherwise the system has great difficulty dealing with their page-consumptive demands.

```
vadvise(VA_NORM);
```

restores default paging replacement behavior after a call to

```
vadvise(VA_ANOM);
```

**BUGS**

The current implementation of **vadvise()** will go away soon, being replaced by a per-page **vadvise()** facility.

**NAME**

**vfork** – spawn new process in a virtual memory efficient way

**SYNOPSIS**

```
#include <vfork.h>
```

```
int vfork()
```

**DESCRIPTION**

**vfork()** can be used to create new processes without fully copying the address space of the old process, which is horrendously inefficient in a paged environment. It is useful when the purpose of **fork(2V)**, would have been to create a new system context for an **execve(2V)**. **vfork()** differs from **fork()** in that the child borrows the parent's memory and thread of control until a call to **execve(2V)**, or an exit (either by a call to **exit(2V)** or abnormally.) The parent process is suspended while the child is using its resources.

**vfork()** returns 0 in the child's context and (later) the process ID (PID) of the child in the parent's context.

**vfork()** can normally be used just like **fork**. It does not work, however, to return while running in the child's context from the procedure which called **vfork()** since the eventual return from **vfork()** would then return to a no longer existent stack frame. Be careful, also, to call **\_exit()** rather than **exit()** if you cannot *execve*, since **exit()** will flush and close standard I/O channels, and thereby mess up the parent processes standard I/O data structures. (Even with **fork()** it is wrong to call **exit()** since buffered data would then be flushed twice.)

On Sun-4 machines, the parent inherits the values of local and incoming argument registers from the child. Since this violates the usual data flow properties of procedure calls, the file **<vfork.h>** must be included in programs that are compiled using global optimization.

**RETURN VALUES**

On success, **vfork()** returns 0 to the child process and returns the process ID of the child process to the parent process. On failure, **vfork()** returns -1 to the parent process, sets **errno** to indicate the error, and no child process is created.

**SEE ALSO**

**execve(2V)**, **exit(2V)**, **fork(2V)**, **ioctl(2)**, **sigvec(2)**, **wait(2V)**

**BUGS**

This system call will be eliminated in a future release. System implementation changes are making the efficiency gain of **vfork()** over **fork(2V)** smaller. The memory sharing semantics of **vfork()** can be obtained through other mechanisms.

To avoid a possible deadlock situation, processes that are children in the middle of a **vfork()** are never sent **SIGTTOU** or **SIGTTIN** signals; rather, output or *ioctls* are allowed and input attempts result in an EOF indication.

**NAME**

**vhangup** – virtually “hangup” the current control terminal

**SYNOPSIS**

**vhangup()**

**DESCRIPTION**

**vhangup()** is used by the initialization process **init(8)** (among others) to arrange that users are given “clean” terminals at login, by revoking access of the previous users’ processes to the terminal. To affect this, **vhangup()** searches the system tables for references to the control terminal of the invoking process, revoking access permissions on each instance of the terminal that it finds. Further attempts to access the terminal by the affected processes will yield I/O errors (EBADF). Finally, a SIGHUP (hangup signal) is sent to the process group of the control terminal.

**SEE ALSO**

**init(8)**

**BUGS**

Access to the control terminal using **/dev/tty** is still possible.

This call should be replaced by an automatic mechanism that takes place on process exit.

**NAME**

wait, wait3, wait4, waitpid, WIFSTOPPED, WIFSIGNALED, WIFEXITED, WEXITSTATUS, WTERMSIG, WSTOPSIG – wait for process to terminate or stop, examine returned status

**SYNOPSIS**

```
#include <sys/wait.h>

int wait(statusp)
int *statusp;

int waitpid(pid, statusp, options)
int pid;
int *statusp;
int options;

#include <sys/time.h>
#include <sys/resource.h>

int wait3(statusp, options, rusage)
int *statusp;
int options;
struct rusage *rusage;

int wait4(pid, statusp, options, rusage)
int pid;
int *statusp;
int options;
struct rusage *rusage;

WIFSTOPPED(status)
int status;

WIFSIGNALED(status)
int status;

WIFEXITED(status)
int status

WEXITSTATUS(status)
int status

WTERMSIG(status)
int status

WSTOPSIG(status)
int status
```

**SYSTEM V SYNOPSIS**

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

pid_t wait(statusp)
int *statusp;

pid_t waitpid(pid, statusp, options)
pid_t pid;
int *statusp;
int options;
```



**DESCRIPTION**

**wait()** delays its caller until a signal is received or one of its child processes terminates or stops due to tracing. If any child has died or stopped due to tracing and this has not been reported using **wait()**, return is immediate, returning the process ID and exit status of one of those children. If that child had died, it is discarded. If there are no children, return is immediate with the value `-1` returned. If there are only running or stopped but reported children, the calling process is blocked.

If *statusp* is not a NULL pointer, then on return from a successful **wait()** call the status of the child process whose process ID is the return value of **wait()** is stored in the location pointed to by *statusp*. It indicates the cause of termination and other information about the terminated process in the following manner:

- If the first byte (the low-order 8 bits) are equal to `0177`, the child process has stopped. The next byte contains the number of the signal that caused the process to stop. See **ptrace(2)** and **sigvec(2)**.
- If the first byte (the low-order 8 bits) are non-zero and are not equal to `0177`, the child process terminated due to a signal. The low-order 7 bits contain the number of the signal that terminated the process. In addition, if the low-order seventh bit (that is, bit `0200`) is set, a "core image" of the process was produced (see **sigvec(2)**).
- Otherwise, the child process terminated due to a call to **exit(2V)**. The next byte contains the low-order 8 bits of the argument that the child process passed to **exit()**.

**waitpid()** behaves identically to **wait()** if *pid* has a value of `-1` and *options* has a value of zero. Otherwise, the behavior of **waitpid()** is modified by the values of *pid* and *options* as follows:

*pid* specifies a set of child processes for which status is requested. **waitpid()** only returns the status of a child process from this set.

- If *pid* is equal to `-1`, status is requested for any child process. In this respect, **waitpid()** is then equivalent to **wait()**.
- If *pid* is greater than zero, it specifies the process ID of a single child process for which status is requested.
- If *pid* is equal to zero, status is requested for any child process whose process group ID is equal to that of the calling process.
- If *pid* is less than `-1`, status is requested for any child process whose process group ID is equal to the absolute value of *pid*.

*options* is constructed from the bitwise inclusive OR of zero or more of the following flags, defined in the header `<sys/wait.h>`:

**WNOHANG**

**waitpid()** does not suspend execution of the calling process if status is not immediately available for one of the child processes specified by *pid*.

**WUNTRACED**

The status of any child processes specified by *pid* that are stopped, and whose status has not yet been reported since they stopped, are also reported to the requesting process.

**wait3()** is an alternate interface that allows both non-blocking status collection and the collection of the status of children stopped by any means. The *status* parameter is defined as above. The *options* parameter is used to indicate the call should not block if there are no processes that have status to report (**WNOHANG**), and/or that children of the current process that are stopped due to a **SIGTTIN**, **SIGTTOU**, **SIGTSTP**, or **SIGSTOP** signal are eligible to have their status reported as well (**WUNTRACED**). A terminated child is discarded after it reports status, and a stopped process will not report its status more than once. If *rusage* is not a NULL pointer, a summary of the resources used by the terminated process and all its children is returned. (This information is currently not available for stopped processes.)

When the **WNOHANG** option is specified and no processes have status to report, **wait3()** returns 0. The **WNOHANG** and **WUNTRACED** options may be combined by ORing the two values.

**wait4()** is another alternate interface. With a *pid* argument of 0, it is equivalent to **wait3()**. If *pid* has a nonzero value, then **wait4()** returns status only for the indicated process ID, but not for any other child processes.

**WIFSTOPPED**, **WIFSIGNALED**, **WIFEXITED**, **WEXITSTATUS**, **WTERMSIG**, and **WSTOPSIG** are macros that take an argument *status*, of type 'int', as returned by **wait()**, **wait3()**, or **wait4()**. **WIFSTOPPED** evaluates to true (1) when the process for which the **wait()** call was made is stopped, or to false (0) otherwise. If **WIFSTOPPED(status)** is non-zero, **WSTOPSIG** evaluates to the number of the signal that caused the child process to stop. **WIFSIGNALED** evaluates to true when the process was terminated with a signal. If **WIFSIGNALED(status)** is non-zero, **WTERMSIG** evaluates to the number of the signal that caused the termination of the child process. **WIFEXITED** evaluates to true when the process exited by using an **exit(2V)** call. If **WIFEXITED(status)** is non-zero, **WEXITSTATUS** evaluates to the low-order byte of the argument that the child process passed to **\_exit()** (see **exit(2V)**) or **exit(3)**, or the value the child process returned from **main()** (see **execve(2V)**).

If the information stored at the location pointed to by *statusp* was stored there by a call to **waitpid()** that specified the **WUNTRACED** flag, exactly one of the macros **WIFEXITED(\*statusp)**, **WIFSIGNALED(\*statusp)**, and **WIFSTOPPED(\*statusp)** will evaluate to a non-zero value. If the information stored at the location pointed to by *statusp* was stored there by a call to **waitpid()** that did *not* specify the **WUNTRACED** flag or by a call to **wait()**, exactly one of the macros **WIFEXITED(\*statusp)** and **WIFSIGNALED(\*statusp)** will evaluate to a non-zero value.

If a parent process terminates without waiting for all of its child processes to terminate, the remaining child processes are assigned the parent process ID of 1, corresponding to **init(8)**.

#### RETURN VALUES

If **wait()** or **waitpid()** returns due to a stopped or terminated child process, the process ID of the child is returned to the calling process. Otherwise, a value of -1 is returned and **errno** is set to indicate the error.

If **wait()** or **waitpid()** return due to the delivery of a signal to the calling process, a value of -1 is returned and **errno** is set to **EINTR**. If **waitpid()** function was invoked with **WNOHANG** set in *options*, it has at least one child process specified by *pid* for which status is not available, and status is not available for any process specified by *pid*, a value of zero is returned. Otherwise, a value of -1 is returned, and **errno** is set to indicate the error.

**wait3()** and **wait4()** return 0 if **WNOHANG** is specified and there are no stopped or exited children, and return the process ID of the child process if they return due to a stopped or terminated child process. Otherwise, they return a value of -1 and set **errno** to indicate the error.

#### ERRORS

**wait()**, **wait3()**, or **wait4()** will fail and return immediately if one or more of the following are true:

<b>ECHILD</b>	The calling process has no existing unwaited-for child processes.
<b>EFAULT</b>	<i>statusp</i> or <i>rusage</i> points to an illegal address.
<b>EINTR</b>	The function was interrupted by a signal. The value of the location pointed to by <i>statusp</i> is undefined.

**waitpid()** may set **errno** to:

<b>ECHILD</b>	The process or process group specified by <i>pid</i> does not exist or is not a child of the calling process.
<b>EINTR</b>	The function was interrupted by a signal. The value of the location pointed to by <i>statusp</i> is undefined.
<b>EINVAL</b>	The value of <i>options</i> is not valid.

**wait()**, **wait3()**, and **wait4()** will terminate prematurely, return `-1`, and set `errno` to: `EINTR` upon the arrival of a signal whose `SV_INTERRUPT` bit in its flags field is set (see **sigvec(2)** and **siginterrupt(3V)**). **signal(3V)**, in the System V compatibility library, sets this bit for any signal it catches.

**SEE ALSO**

**exit(2V)**, **fork(2V)**, **getrusage(2)**, **ptrace(2)**, **sigvec(2)**, **pause(3V)**, **siginterrupt(3V)**, **signal(3V)**, **times(3V)**

**NOTES**

If a parent process terminates without waiting on its children, the initialization process (process ID = 1) inherits the children.

**wait()**, **wait3()**, and **wait4()** are automatically restarted when a process receives a signal while awaiting termination of a child process, unless the `SV_INTERRUPT` bit is set in the flags for that signal.

Previous SunOS releases used **union wait \*statusp** and **union wait status** in place of **int \*statusp** and **intstatus**. The union contained a member `w_status` that could be treated in the same way as `status`.

Other members of the `wait` union could be used to extract this information more conveniently:

- If the `w_stopval` member had the value `WSTOPPED`, the child process had stopped; the value of the `w_stopsig` member was the signal that stopped the process.
- If the `w_termsig` member was non-zero, the child process terminated due to a signal; the value of the `w_termsig` member was the number of the signal that terminated the process. If the `w_coredump` member was non-zero, a core dump was produced.
- Otherwise, the child process terminated due to a call to `exit()`. The value of the `w_retcode` member was the low-order 8 bits of the argument that the child process passed to `exit()`.

**union wait** is obsolete in light of the new specifications provided by *IEEE Std 1003.1-1988* and endorsed by *SVID89* and *XPG3*. SunOS Release 4.1 supports **union wait** for backward compatibility, but it will disappear in a future release.

## NAME

write, writev – write output

## SYNOPSIS

```
int write(fd, buf, nbyte)
int fd;
char *buf;
int nbyte;

#include <sys/types.h>
#include <sys/uio.h>

int writev(fd, iov, iovcnt)
int fd;
struct iovec *iov;
int iovcnt;
```

## SYSTEM V SYNOPSIS

```
int write(fd, buf, nbyte)
int fd;
char *buf;
unsigned nbyte;
```

## DESCRIPTION

**write()** attempts to write *nbyte* bytes of data to the object referenced by the descriptor *fd* from the buffer pointed to by *buf*. **writev()** performs the same action, but gathers the output data from the *iovcnt* buffers specified by the members of the *iov* array: *iov[0]*, *iov[1]*, ..., *iov[iovcnt - 1]*. If *nbyte* is zero, **write()** takes no action and returns 0. **writev()**, however, returns -1 and sets the global variable *errno* (see ERRORS below).

For **writev()**, the *iovec* structure is defined as

```
struct iovec {
    caddr_t iov_base;
    int     iov_len;
};
```

Each *iovec* entry specifies the base address and length of an area in memory from which data should be written. **writev()** always writes a complete area before proceeding to the next.

On objects capable of seeking, the **write()** starts at a position given by the seek pointer associated with *fd*, (see **lseek(2V)**). Upon return from **write()**, the seek pointer is incremented by the number of bytes actually written.

Objects that are not capable of seeking always write from the current position. The value of the seek pointer associated with such an object is undefined.

If the **O\_APPEND** flag of the file status flags is set, the seek pointer is set to the end of the file prior to each write.

If the process calling **write()** or **writev()** receives a signal before any data are written, the system call is restarted, unless the process explicitly set the signal to interrupt the call using **sigvec()** or **sigaction()** (see the discussions of **SV\_INTERRUPT** on **sigvec(2)** and **SA\_INTERRUPT** on **sigaction(3V)**). If **write()** or **writev()** is interrupted by a signal after successfully writing some data, it returns the number of bytes written.

For regular files, if the **O\_SYNC** flag of the file status flags is set, **write()** does not return until both the file data and file status have been physically updated. This function is for special applications that require extra reliability at the cost of performance. For block special files, if **O\_SYNC** is set, the **write()** does not return until the data has been physically updated.

If the real user is not the super-user, then `write()` clears the set-user-id bit on a file. This prevents penetration of system security by a user who "captures" a writable set-user-id file owned by the super-user.

For STREAMS (see `intro(2)`) files, the operation of `write()` and `writev()` are determined by the values of the minimum and maximum packet sizes accepted by the *stream*. These values are contained in the topmost *stream* module. Unless the user pushes (see `I_PUSH` in `streamio(4)`) the topmost module, these values can not be set or tested from user level. If the total number of bytes to be written falls within the packet size range, that many bytes are written. If the total number of bytes to be written does not fall within the range and the minimum packet size value is zero, `write()` and `writev()` break the data to be written into maximum packet size segments prior to sending the data downstream (the last segment may contain less than the maximum packet size). If the total number of bytes to be written does not fall within the range and the minimum value is non-zero, `write()` and `writev()` fail and set `errno` to `ERANGE`. Writing a zero-length buffer (the total number of bytes to be written is zero) sends zero bytes with zero returned.

When a descriptor or the object it refers to is marked for non-blocking I/O, and the descriptor refers to an object subject to flow control, such as a socket, a pipe (or FIFO), or a *stream*, `write()` and `writev()` may write fewer bytes than requested; the return value must be noted, and the remainder of the operation should be retried when possible. If such an object's buffers are full, so that it cannot accept any data, then:

- If the object to which the descriptor refers is marked for non-blocking I/O using the `FIONBIO` request to `ioctl(2)`, or by using `fcntl(2V)` to set the `FNDELAY` or `O_NDELAY` flag (defined in `<sys/fcntl.h>`), `write()` returns `-1` and sets `errno` to `EWOULDBLOCK`.

Upon successful completion, `write()` marks for update the `st_ctime` and `st_mtime` fields of the file.

#### SYSTEM V DESCRIPTION

`write()` and `writev()` behave as described above, except:

When a descriptor or the object it refers to is marked for non-blocking I/O, and the descriptor refers to an object subject to flow control, such as a socket, a pipe (or FIFO), or a *stream*, `write()` and `writev()` may write fewer bytes than requested; the return value must be noted, and the remainder of the operation should be retried when possible. If such an object's buffers are full, so that it cannot accept any data, then:

- If the descriptor is marked for non-blocking I/O by using `fcntl()` to set the `FNDELAY` or `O_NDELAY` flag (defined in `<sys/fcntl.h>`), and does not refer to a *stream*, the `write()` returns 0. If the descriptor is marked for non-blocking I/O, and refers to a *stream*, `write()` returns `-1` and sets `errno` to `EAGAIN`.
- If the descriptor is marked for non-blocking I/O using `fcntl()` to set the `FNONBLOCK` or `O_NONBLOCK` flag (defined in `<sys/fcntl.h>`), `write()` requests for `{PIPE_BUF}` (see `pathconf(2V)`) or fewer bytes either succeed completely and return `nbyte`, or return `-1` and set `errno` to `EAGAIN`. A `write()` request for greater than `{PIPE_BUF}` bytes either transfers what it can and returns the number of bytes written, or transfers no data and returns `-1` and sets `errno` to `EAGAIN`. If a `write()` request is greater than `{PIPE_BUF}` bytes and all data previously written to the pipe has been read, `write()` transfers at least `{PIPE_BUF}` bytes.

#### RETURN VALUES

`write()` and `writev()` return the number of bytes actually written on success. On failure, they return `-1` and set `errno` to indicate the error.

#### ERRORS

`write()` and `writev()` fail and the seek pointer remains unchanged if one or more of the following are true:

<code>EBADF</code>	<i>fd</i> is not a valid descriptor open for writing.
<code>EDQUOT</code>	The user's quota of disk blocks on the file system containing the file has been exhausted.
<code>EFAULT</code>	Part of <i>iov</i> or data to be written to the file points outside the process's allocated address space.

EFBIG	An attempt was made to write a file that exceeds the process's file size limit or the maximum file size.
EINTR	The process performing a write received a signal before any data were written, and the signal was set to interrupt the system call.
EINVAL	The <i>stream</i> is linked below a multiplexor. The seek pointer associated with <i>fd</i> was negative.
EIO	An I/O error occurred while reading from or writing to the file system. The process is in a background process group and is attempting to write to its controlling terminal, TOSTOP is set, the process is neither ignoring nor blocking SIGTTOU, and the process group of the process is orphaned.
ENOSPC	There is no free space remaining on the file system containing the file.
ENXIO	A hangup occurred on the <i>stream</i> being written to.
EPIPE	An attempt is made to write to a pipe that is not open for reading by any process (or to a socket of type SOCK_STREAM that is connected to a peer socket.) Note: an attempted write of this kind also causes you to receive a SIGPIPE signal from the kernel. If you've not made a special provision to catch or ignore this signal, then your process dies.
ERANGE	<i>fd</i> refers to a <i>stream</i> , the total number of bytes to be written is outside the minimum and maximum write range, and the minimum value is non-zero.
EWOLDBLOCK	The file was marked for non-blocking I/O, and no data could be written immediately.

In addition to the above, `writev()` may set `errno` to:

EINVAL	<i>iovcnt</i> was less than or equal to 0, or greater than 16. One of the <i>iov_len</i> values in the <i>iov</i> array was negative. The sum of the <i>iov_len</i> values in the <i>iov</i> array overflowed a 32-bit integer.
--------	---

A write to a STREAMS file can fail if an error message has been received at the stream head. In this case, `errno` is set to the value included in the error message.

#### SYSTEM V ERRORS

`write()` fails and sets `errno` as described above, except:

EAGAIN	The descriptor referred to a <i>stream</i> , was marked for non-blocking I/O, and no data could be written immediately. The O_NONBLOCK flag is set for the file descriptor and <code>write()</code> would block.
--------	---

#### SEE ALSO

`dup(2V)`, `fcntl(2V)`, `intro(2)`, `ioctl(2)`, `lseek(2V)`, `open(2V)`, `pipe(2V)`, `select(2)`, `sigvec(2)`, `signal(3V)`

**NAME**

intro – introduction to user-level library functions

**DESCRIPTION**

Section 3 describes user-level library routines. In this release, most user-library routines are listed in alphabetical order regardless of their subsection headings. (This eliminates having to search through several subsections of the manual.) However, due to their special-purpose nature, the routines from the following libraries are broken out into the indicated subsections:

- The Lightweight Processes Library, in subsection 3L.
- The Mathematical Library, in subsection 3M.
- The RPC Services Library, in subsection 3R.

A 3V section number means one or more of the following:

- The man page documents System V behavior only.
- The man page documents default SunOS behavior, and System V behavior as it differs from the default behavior. These System V differences are presented under **SYSTEM V** section headers.
- The man page documents behavior compliant with *IEEE Std 1003.1-1988* (POSIX.1).

The System V Library was formerly documented in a separate manual section. These man pages have been merged into the main portion of section 3. These man pages describe functions that may differ from the default SunOS functions. To use them, compile programs with `/usr/5bin/cc` instead of `/usr/bin/cc`.

Section 3 also documents the library interfaces for *X/Open Portability Guide, Issue 2* (XPG2) compatibility. Where these interfaces differ from the System V versions, the differences are noted. To use the XPG2 compatibility library interfaces, compile programs with `/usr/xpg2bin/cc`.

The libraries provide many different “standard” environments. These environments (including two that are not yet fully supported) are described on `ansic(7V)`, `bsd(7)`, `posix(7V)`, `sunos(7)`, `svidii(7V)`, `svidiii(7V)`, and `xopen(7V)`.

The main C library, `/usr/lib/libc.a`, contains many of the functions described in this section, along with entry points for the system calls described in Section 2. This library also includes the Internet networking routines listed under the 3N subsection heading, and routines provided for compatibility with other UNIX operating systems, listed under 3C. Functions associated with the “standard I/O library” are listed under 3S.

User-level routines for access to data structures within the kernel and other processes are listed under 3K. To use these functions, compile programs with the `-lkvm` option for the C compiler, `cc(1V)`.

Math library functions are listed under 3M. To use them, compile programs with the `-lm` `cc(1V)` option.

Various specialized libraries, the routines they contain, and the compiler options needed to link with them, are listed under 3X.

**FILES**

<code>/usr/lib/libc.a</code>	C Library (2, 3, 3N and 3C)
<code>/usr/lib/lib*.a</code>	other “standard” C libraries
<code>/usr/lib/lib*.a</code>	special-purpose C libraries
<code>/usr/5bin/cc</code>	

**SEE ALSO**

`cc(1V)`, `ld(1)`, `nm(1)`, `intro(2)`

## LIST OF LIBRARY FUNCTIONS

Name	Appears on Page	Description
<b>a64l</b>	<b>a64l(3)</b>	convert between long integer and base-64 ASCII string
<b>abort</b>	<b>abort(3)</b>	generate a fault
<b>abs</b>	<b>abs(3)</b>	integer absolute value
<b>addexportent</b>	<b>exportent(3)</b>	get exported file system information
<b>addmntent</b>	<b>getmntent(3)</b>	get file system descriptor file entry
<b>aiocancel</b>	<b>aiocancel(3)</b>	cancel an asynchronous operation
<b>aioread</b>	<b>aioread(3)</b>	asynchronous I/O operations
<b>aiowait</b>	<b>aiowait(3)</b>	wait for completion of asynchronous I/O operation
<b>aiowrite</b>	<b>aioread(3)</b>	asynchronous I/O operations
<b>alarm</b>	<b>alarm(3V)</b>	schedule signal after specified time
<b>alloca</b>	<b>malloc(3V)</b>	memory allocator
<b>alphasort</b>	<b>scandir(3)</b>	scan a directory
<b>arc</b>	<b>plot(3X)</b>	graphics interface
<b>asctime</b>	<b>ctime(3V)</b>	convert date and time
<b>assert</b>	<b>assert(3V)</b>	program verification
<b>atof</b>	<b>strtod(3)</b>	convert string to double-precision number
<b>atoi</b>	<b>strtol(3)</b>	convert string to integer
<b>atol</b>	<b>strtol(3)</b>	convert string to integer
<b>audit_args</b>	<b>audit_args(3)</b>	produce text audit message
<b>audit_text</b>	<b>audit_args(3)</b>	produce text audit message
<b>auth_destroy</b>	<b>rpc_clnt_auth(3N)</b>	library routines for client side RPC authentication
<b>authdes_create</b>	<b>secure_rpc(3N)</b>	library routines for secure remote procedure calls
<b>authdes_getcred</b>	<b>secure_rpc(3N)</b>	library routines for secure remote procedure calls
<b>authnone_create</b>	<b>rpc_clnt_auth(3N)</b>	library routines for client side RPC authentication
<b>authunix_create</b>	<b>rpc_clnt_auth(3N)</b>	library routines for client side RPC authentication
<b>authunix_create_default</b>	<b>rpc_clnt_auth(3N)</b>	library routines for client side RPC authentication
<b>bcmp</b>	<b>bstring(3)</b>	bit and byte string operations
<b>bcopy</b>	<b>bstring(3)</b>	bit and byte string operations
<b>bindresvport</b>	<b>bindresvport(3N)</b>	bind a socket to a privileged IP port
<b>bsearch</b>	<b>bsearch(3)</b>	binary search a sorted table
<b>bstring</b>	<b>bstring(3)</b>	bit and byte string operations
<b>byteorder</b>	<b>byteorder(3N)</b>	convert values between host and network byte order
<b>bzero</b>	<b>bstring(3)</b>	bit and byte string operations
<b>calloc</b>	<b>malloc(3V)</b>	memory allocator
<b>callrpc</b>	<b>rpc_clnt_calls(3N)</b>	library routines for client side calls
<b>catclose</b>	<b>catopen(3C)</b>	open/close a message catalog
<b>catgetmsg</b>	<b>catgets(3C)</b>	get message from a message catalog
<b>catgets</b>	<b>catgets(3C)</b>	get message from a message catalog
<b>catopen</b>	<b>catopen(3C)</b>	open/close a message catalog
<b>cbc_crypt</b>	<b>des_crypt(3)</b>	fast DES encryption
<b>cfgetispeed</b>	<b>termios(3V)</b>	terminal control functions
<b>cfgetospeed</b>	<b>termios(3V)</b>	terminal control functions
<b>cfree</b>	<b>malloc(3V)</b>	memory allocator
<b>cfsetispeed</b>	<b>termios(3V)</b>	terminal control functions
<b>cfsetospeed</b>	<b>termios(3V)</b>	terminal control functions
<b>circle</b>	<b>plot(3X)</b>	graphics interface
<b>clearerr</b>	<b>ferror(3V)</b>	stream status inquiries
<b>clnt_broadcast</b>	<b>rpc_clnt_calls(3N)</b>	library routines for client side calls
<b>clnt_call</b>	<b>rpc_clnt_calls(3N)</b>	library routines for client side calls
<b>clnt_control</b>	<b>rpc_clnt_create(3N)</b>	library routines creating and manipulating CLIENT handles



<b>clnt_create</b>	<b>rpc_clnt_create(3N)</b>	library routines creating and manipulating CLIENT handles
<b>clnt_create_vers</b>	<b>rpc_clnt_create(3N)</b>	library routines creating and manipulating CLIENT handles
<b>clnt_destroy</b>	<b>rpc_clnt_create(3N)</b>	library routines creating and manipulating CLIENT handles
<b>clnt_freeres</b>	<b>rpc_clnt_calls(3N)</b>	library routines for client side calls
<b>clnt_geterr</b>	<b>rpc_clnt_calls(3N)</b>	library routines for client side calls
<b>clnt_pcreateerror</b>	<b>rpc_clnt_create(3N)</b>	library routines creating and manipulating CLIENT handles
<b>clnt_perrno</b>	<b>rpc_clnt_calls(3N)</b>	library routines for client side calls
<b>clnt_perror</b>	<b>rpc_clnt_calls(3N)</b>	library routines for client side calls
<b>clnt_screateerror</b>	<b>rpc_clnt_create(3N)</b>	library routines creating and manipulating CLIENT handles
<b>clnt_sperrno</b>	<b>rpc_clnt_calls(3N)</b>	library routines for client side calls
<b>clnt_sperror</b>	<b>rpc_clnt_calls(3N)</b>	library routines for client side calls
<b>clntraw_create</b>	<b>rpc_clnt_create(3N)</b>	library routines creating and manipulating CLIENT handles
<b>clnttcp_create</b>	<b>rpc_clnt_create(3N)</b>	library routines creating and manipulating CLIENT handles
<b>clntudp_bufcreate</b>	<b>rpc_clnt_create(3N)</b>	library routines creating and manipulating CLIENT handles
<b>clock</b>	<b>clock(3C)</b>	report CPU time used
<b>closedir</b>	<b>directory(3V)</b>	directory operations
<b>closelog</b>	<b>syslog(3)</b>	control system log
<b>closepl</b>	<b>plot(3X)</b>	graphics interface
<b>cont</b>	<b>plot(3X)</b>	graphics interface
<b>conv</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>crypt</b>	<b>crypt(3)</b>	password and data encryption
<b>ctermid</b>	<b>ctermid(3V)</b>	generate filename for terminal
<b>ctime</b>	<b>ctime(3V)</b>	convert date and time
<b>ctype</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>curses</b>	<b>curses(3V)</b>	System V terminal screen handling and optimization package
<b>cuserid</b>	<b>cuserid(3V)</b>	get character login name of the user
<b>dbm</b>	<b>dbm(3X)</b>	data base subroutines
<b>dbm_clearerr</b>	<b>ndbm(3)</b>	data base subroutines
<b>dbm_close</b>	<b>ndbm(3)</b>	data base subroutines
<b>dbm_delete</b>	<b>ndbm(3)</b>	data base subroutines
<b>dbm_error</b>	<b>ndbm(3)</b>	data base subroutines
<b>dbm_fetch</b>	<b>ndbm(3)</b>	data base subroutines
<b>dbm_firstkey</b>	<b>ndbm(3)</b>	data base subroutines
<b>dbm_nextkey</b>	<b>ndbm(3)</b>	data base subroutines
<b>dbm_open</b>	<b>ndbm(3)</b>	data base subroutines
<b>dbm_store</b>	<b>ndbm(3)</b>	data base subroutines
<b>dbmclose</b>	<b>dbm(3X)</b>	data base subroutines
<b>dbminit</b>	<b>dbm(3X)</b>	data base subroutines
<b>decimal_to_double</b>	<b>decimal_to_floating(3)</b>	convert decimal record to floating-point value
<b>decimal_to_extended</b>	<b>decimal_to_floating(3)</b>	convert decimal record to floating-point value
<b>decimal_to_single</b>	<b>decimal_to_floating(3)</b>	convert decimal record to floating-point value
<b>delete</b>	<b>dbm(3X)</b>	data base subroutines
<b>des_crypt</b>	<b>des_crypt(3)</b>	fast DES encryption
<b>des_setparity</b>	<b>des_crypt(3)</b>	fast DES encryption
<b>directory</b>	<b>directory(3V)</b>	directory operations
<b>dlclose</b>	<b>dlopen(3X)</b>	simple programmatic interface to the dynamic linker
<b>dlerror</b>	<b>dlopen(3X)</b>	simple programmatic interface to the dynamic linker
<b>dlopen</b>	<b>dlopen(3X)</b>	simple programmatic interface to the dynamic linker
<b>dlsym</b>	<b>dlopen(3X)</b>	simple programmatic interface to the dynamic linker
<b>dn_comp</b>	<b>resolver(3)</b>	resolver routines
<b>dn_expand</b>	<b>resolver(3)</b>	resolver routines
<b>double_to_decimal</b>	<b>floating_to_decimal(3)</b>	convert floating-point value to decimal record
<b>drand48</b>	<b>drand48(3)</b>	generate uniformly distributed pseudo-random numbers

<b>dysize</b>	<b>ctime(3V)</b>	convert date and time
<b>ecb_crypt</b>	<b>des_crypt(3)</b>	fast DES encryption
<b>econvert</b>	<b>econvert(3)</b>	output conversion
<b>ecvt</b>	<b>econvert(3)</b>	output conversion
<b>edata</b>	<b>end(3)</b>	last locations in program
<b>encrypt</b>	<b>crypt(3)</b>	password and data encryption
<b>end</b>	<b>end(3)</b>	last locations in program
<b>endac</b>	<b>getacinfo(3)</b>	get audit control file information
<b>endexportent</b>	<b>exportent(3)</b>	get exported file system information
<b>endsent</b>	<b>getsent(3)</b>	get file system descriptor file entry
<b>endgraent</b>	<b>getgraent(3)</b>	get group adjunct file entry
<b>endgrent</b>	<b>getgrent(3V)</b>	get group file entry
<b>endhostent</b>	<b>gethostent(3N)</b>	get network host entry
<b>endmntent</b>	<b>getmntent(3)</b>	get file system descriptor file entry
<b>endnetent</b>	<b>getnetent(3N)</b>	get network entry
<b>endnetgrent</b>	<b>getnetgrent(3N)</b>	get network group entry
<b>endprotoent</b>	<b>getprotoent(3N)</b>	get protocol entry
<b>endpwaent</b>	<b>getpwaent(3)</b>	get password adjunct file entry
<b>endpwent</b>	<b>getpwent(3V)</b>	get password file entry
<b>endrpcent</b>	<b>getrpcent(3N)</b>	get RPC entry
<b>endservent</b>	<b>getservent(3N)</b>	get service entry
<b>endttyent</b>	<b>getttyent(3)</b>	get ttytab file entry
<b>endusershell</b>	<b>getusershell(3)</b>	get legal user shells
<b>erand48</b>	<b>drand48(3)</b>	generate uniformly distributed pseudo-random numbers
<b>erase</b>	<b>plot(3X)</b>	graphics interface
<b>errno</b>	<b>perror(3)</b>	system error messages
<b>etext</b>	<b>end(3)</b>	last locations in program
<b>ether_aton</b>	<b>ethers(3N)</b>	Ethernet address mapping operations
<b>ether_hostton</b>	<b>ethers(3N)</b>	Ethernet address mapping operations
<b>ether_line</b>	<b>ethers(3N)</b>	Ethernet address mapping operations
<b>ether_ntoa</b>	<b>ethers(3N)</b>	Ethernet address mapping operations
<b>ether_ntohost</b>	<b>ethers(3N)</b>	Ethernet address mapping operations
<b>ethers</b>	<b>ethers(3N)</b>	Ethernet address mapping operations
<b>execl</b>	<b>execl(3V)</b>	execute a file
<b>execle</b>	<b>execl(3V)</b>	execute a file
<b>execlp</b>	<b>execl(3V)</b>	execute a file
<b>execv</b>	<b>execl(3V)</b>	execute a file
<b>execvp</b>	<b>execl(3V)</b>	execute a file
<b>exit</b>	<b>exit(3)</b>	terminate a process after performing cleanup
<b>exportent</b>	<b>exportent(3)</b>	get exported file system information
<b>extended_to_decimal</b>	<b>floating_to_decimal(3)</b>	convert floating-point value to decimal record
<b>fclose</b>	<b>fclose(3V)</b>	close or flush a stream
<b>fconvert</b>	<b>econvert(3)</b>	output conversion
<b>fcvt</b>	<b>econvert(3)</b>	output conversion
<b>fdopen</b>	<b>fopen(3V)</b>	open a stream
<b>feof</b>	<b>ferror(3V)</b>	stream status inquiries
<b>ferror</b>	<b>ferror(3V)</b>	stream status inquiries
<b>fetch</b>	<b>dbm(3X)</b>	data base subroutines
<b>fflush</b>	<b>fclose(3V)</b>	close or flush a stream
<b>ffs</b>	<b>bstring(3)</b>	bit and byte string operations
<b>fgetc</b>	<b>getc(3V)</b>	get character or integer from stream
<b>fgetgraent</b>	<b>getgraent(3)</b>	get group adjunct file entry
<b>fgetgrent</b>	<b>getgrent(3V)</b>	get group file entry

<b>fgetpwaent</b>	<b>getpwaent(3)</b>	get password adjunct file entry
<b>fgetpwent</b>	<b>getpwent(3V)</b>	get password file entry
<b>fgets</b>	<b>gets(3S)</b>	get a string from a stream
<b>file_to_decimal</b>	<b>string_to_decimal(3)</b>	parse characters into decimal record
<b>fileno</b>	<b>ferror(3V)</b>	stream status inquiries
<b>firstkey</b>	<b>dbm(3X)</b>	data base subroutines
<b>floatingpoint</b>	<b>floatingpoint(3)</b>	IEEE floating point definitions
<b>fopen</b>	<b>fopen(3V)</b>	open a stream
<b>fprintf</b>	<b>printf(3V)</b>	formatted output conversion
<b>fputc</b>	<b>putc(3S)</b>	put character or word on a stream
<b>fputs</b>	<b>puts(3S)</b>	put a string on a stream
<b>fread</b>	<b>fread(3S)</b>	buffered binary input/output
<b>free</b>	<b>malloc(3V)</b>	memory allocator
<b>freopen</b>	<b>fopen(3V)</b>	open a stream
<b>fscanf</b>	<b>scanf(3V)</b>	formatted input conversion
<b>fseek</b>	<b>fseek(3S)</b>	reposition a stream
<b>ftell</b>	<b>fseek(3S)</b>	reposition a stream
<b>ftime</b>	<b>time(3V)</b>	get date and time
<b>ftok</b>	<b>ftok(3)</b>	standard interprocess communication package
<b>ftw</b>	<b>ftw(3)</b>	walk a file tree
<b>func_to_decimal</b>	<b>string_to_decimal(3)</b>	parse characters into decimal record
<b>fwrite</b>	<b>fread(3S)</b>	buffered binary input/output
<b>gcd</b>	<b>mp(3X)</b>	multiple precision integer arithmetic
<b>gconvert</b>	<b>econvert(3)</b>	output conversion
<b>gcvrt</b>	<b>econvert(3)</b>	output conversion
<b>get_myaddress</b>	<b>secure_rpc(3N)</b>	library routines for secure remote procedure calls
<b>getacdir</b>	<b>getacinfo(3)</b>	get audit control file information
<b>getacflg</b>	<b>getacinfo(3)</b>	get audit control file information
<b>getacinfo</b>	<b>getacinfo(3)</b>	get audit control file information
<b>getacmin</b>	<b>getacinfo(3)</b>	get audit control file information
<b>getauditflagsbin</b>	<b>getauditflags(3)</b>	convert audit flag specifications
<b>getauditflagschar</b>	<b>getauditflags(3)</b>	convert audit flag specifications
<b>getc</b>	<b>getc(3V)</b>	get character or integer from stream
<b>getchar</b>	<b>getc(3V)</b>	get character or integer from stream
<b>getcwd</b>	<b>getcwd(3V)</b>	get pathname of current working directory
<b>getenv</b>	<b>getenv(3V)</b>	return value for environment name
<b>getexportent</b>	<b>exportent(3)</b>	get exported file system information
<b>getexportopt</b>	<b>exportent(3)</b>	get exported file system information
<b>getfauditflags</b>	<b>getfauditflags(3)</b>	generates the process audit state
<b>getfsent</b>	<b>getfsent(3)</b>	get file system descriptor file entry
<b>getfsfile</b>	<b>getfsent(3)</b>	get file system descriptor file entry
<b>getfsspec</b>	<b>getfsent(3)</b>	get file system descriptor file entry
<b>getfstype</b>	<b>getfsent(3)</b>	get file system descriptor file entry
<b>getgraent</b>	<b>getgraent(3)</b>	get group adjunct file entry
<b>getgranam</b>	<b>getgraent(3)</b>	get group adjunct file entry
<b>getgrent</b>	<b>getgrent(3V)</b>	get group file entry
<b>getgrgid</b>	<b>getgrent(3V)</b>	get group file entry
<b>getgrnam</b>	<b>getgrent(3V)</b>	get group file entry
<b>gethostbyaddr</b>	<b>gethostent(3N)</b>	get network host entry
<b>gethostbyname</b>	<b>gethostent(3N)</b>	get network host entry
<b>gethostent</b>	<b>gethostent(3N)</b>	get network host entry
<b>getlogin</b>	<b>getlogin(3V)</b>	get login name
<b>getmntent</b>	<b>getmntent(3)</b>	get file system descriptor file entry

<b>getnetbyaddr</b>	<b>getnetent(3N)</b>	get network entry
<b>getnetbyname</b>	<b>getnetent(3N)</b>	get network entry
<b>getnetent</b>	<b>getnetent(3N)</b>	get network entry
<b>getnetgrent</b>	<b>getnetgrent(3N)</b>	get network group entry
<b>getnetname</b>	<b>secure_rpc(3N)</b>	library routines for secure remote procedure calls
<b>getopt</b>	<b>getopt(3)</b>	get option letter from argument vector
<b>getpass</b>	<b>getpass(3V)</b>	read a password
<b>getprotobyname</b>	<b>getprotoent(3N)</b>	get protocol entry
<b>getprotobynumber</b>	<b>getprotoent(3N)</b>	get protocol entry
<b>getprotoent</b>	<b>getprotoent(3N)</b>	get protocol entry
<b>getpublickey</b>	<b>publickey(3R)</b>	get public or secret key
<b>getpw</b>	<b>getpw(3)</b>	get name from uid
<b>getpwaent</b>	<b>getpwaent(3)</b>	get password adjunct file entry
<b>getpwanam</b>	<b>getpwaent(3)</b>	get password adjunct file entry
<b>getpwent</b>	<b>getpwent(3V)</b>	get password file entry
<b>getpwnam</b>	<b>getpwent(3V)</b>	get password file entry
<b>getpwuid</b>	<b>getpwent(3V)</b>	get password file entry
<b>getrpcbyname</b>	<b>getrpcent(3N)</b>	get RPC entry
<b>getrpcbynumber</b>	<b>getrpcent(3N)</b>	get RPC entry
<b>getrpcent</b>	<b>getrpcent(3N)</b>	get RPC entry
<b>gets</b>	<b>gets(3S)</b>	get a string from a stream
<b>getsecretkey</b>	<b>publickey(3R)</b>	get public or secret key
<b>getservbyname</b>	<b>getservent(3N)</b>	get service entry
<b>getservbyport</b>	<b>getservent(3N)</b>	get service entry
<b>getservent</b>	<b>getservent(3N)</b>	get service entry
<b>getsubopt</b>	<b>getsubopt(3)</b>	parse sub options from a string.
<b>gettext</b>	<b>gettext(3)</b>	retrieve a message string, get and set text domain
<b>gettyent</b>	<b>gettyent(3)</b>	get ttytab file entry
<b>gettyenam</b>	<b>gettyent(3)</b>	get ttytab file entry
<b>getusershell</b>	<b>getusershell(3)</b>	get legal user shells
<b>getw</b>	<b>getc(3V)</b>	get character or integer from stream
<b>getwd</b>	<b>getwd(3)</b>	get current working directory pathname
<b>gmtime</b>	<b>ctime(3V)</b>	convert date and time
<b>grpauth</b>	<b>pwdauth(3)</b>	password authentication routines
<b>gsignal</b>	<b>ssignal(3)</b>	software signals
<b>gtty</b>	<b>stty(3C)</b>	set and get terminal state
<b>hasmntopt</b>	<b>getmntent(3)</b>	get file system descriptor file entry
<b>hcreate</b>	<b>hsearch(3)</b>	manage hash search tables
<b>hdestroy</b>	<b>hsearch(3)</b>	manage hash search tables
<b>host2netname</b>	<b>secure_rpc(3N)</b>	library routines for secure remote procedure calls
<b>hsearch</b>	<b>hsearch(3)</b>	manage hash search tables
<b>htonl</b>	<b>byteorder(3N)</b>	convert values between host and network byte order
<b>htons</b>	<b>byteorder(3N)</b>	convert values between host and network byte order
<b>index</b>	<b>string(3)</b>	string operations
<b>inet</b>	<b>inet(3N)</b>	Internet address manipulation
<b>inet_addr</b>	<b>inet(3N)</b>	Internet address manipulation
<b>inet_lnaof</b>	<b>inet(3N)</b>	Internet address manipulation
<b>inet_makeaddr</b>	<b>inet(3N)</b>	Internet address manipulation
<b>inet_netof</b>	<b>inet(3N)</b>	Internet address manipulation
<b>inet_network</b>	<b>inet(3N)</b>	Internet address manipulation
<b>inet_ntoa</b>	<b>inet(3N)</b>	Internet address manipulation
<b>initgroups</b>	<b>initgroups(3)</b>	initialize supplementary group IDs
<b>initstate</b>	<b>random(3)</b>	better random number generator

<b>innetgr</b>	<b>getnetgrent(3N)</b>	get network group entry
<b>insque</b>	<b>insque(3)</b>	insert/remove element from a queue
<b>isalnum</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>isalpha</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>isascii</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>isatty</b>	<b>ttyname(3V)</b>	find name of a terminal
<b>iscntrl</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>isdigit</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>isgraph</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>islower</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>isprint</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>ispunct</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>issecure</b>	<b>issecure(3)</b>	character classification and conversion macros and functions
<b>isspace</b>	<b>ctype(3V)</b>	indicates whether system is running secure
<b>isupper</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>isxdigit</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>itom</b>	<b>mp(3X)</b>	character classification and conversion macros and functions
<b>jrand48</b>	<b>drand48(3)</b>	multiple precision integer arithmetic
<b>key_decryptsession</b>	<b>secure_rpc(3N)</b>	generate uniformly distributed pseudo-random numbers
<b>key_encryptsession</b>	<b>secure_rpc(3N)</b>	library routines for secure remote procedure calls
<b>key_gendes</b>	<b>secure_rpc(3N)</b>	library routines for secure remote procedure calls
<b>key_setsecret</b>	<b>secure_rpc(3N)</b>	library routines for secure remote procedure calls
<b>kvm_close</b>	<b>kvm_open(3K)</b>	library routines for secure remote procedure calls
<b>kvm_getcmd</b>	<b>kvm_getu(3K)</b>	specify a kernel to examine
<b>kvm_getproc</b>	<b>kvm_nextproc(3K)</b>	get the u-area or invocation arguments for a process
<b>kvm_getu</b>	<b>kvm_getu(3K)</b>	read system process structures
<b>kvm_nextproc</b>	<b>kvm_nextproc(3K)</b>	get the u-area or invocation arguments for a process
<b>kvm_nlist</b>	<b>kvm_nlist(3K)</b>	read system process structures
<b>kvm_open</b>	<b>kvm_open(3K)</b>	get entries from kernel symbol table
<b>kvm_read</b>	<b>kvm_read(3K)</b>	specify a kernel to examine
<b>kvm_setproc</b>	<b>kvm_nextproc(3K)</b>	copy data to or from a kernel image or running system
<b>kvm_write</b>	<b>kvm_read(3K)</b>	read system process structures
<b>l3tol</b>	<b>l3tol(3C)</b>	copy data to or from a kernel image or running system
<b>l64a</b>	<b>a64l(3)</b>	convert between 3-byte integers and long integers
<b>label</b>	<b>plot(3X)</b>	convert between long integer and base-64 ASCII string
<b>lcong48</b>	<b>drand48(3)</b>	graphics interface
<b>ldaclose</b>	<b>ldclose(3X)</b>	generate uniformly distributed pseudo-random numbers
<b>ldahread</b>	<b>ldahread(3X)</b>	close a COFF file
<b>ldaopen</b>	<b>ldopen(3X)</b>	read the archive header of a member of a COFF archive file
<b>ldclose</b>	<b>ldclose(3X)</b>	open a COFF file for reading
<b>ldfcn</b>	<b>ldfcn(3)</b>	close a COFF file
<b>ldfhread</b>	<b>ldfhread(3X)</b>	common object file access routines
<b>ldgetname</b>	<b>ldgetname(3X)</b>	read the file header of a COFF file
<b>ldlinit</b>	<b>ldlread(3X)</b>	retrieve symbol name for COFF file symbol table entry
<b>ldlitem</b>	<b>ldlread(3X)</b>	manipulate line number entries of a COFF file function
<b>ldlread</b>	<b>ldlread(3X)</b>	manipulate line number entries of a COFF file function
<b>ldlseek</b>	<b>ldlseek(3X)</b>	manipulate line number entries of a COFF file function
<b>ldlnseek</b>	<b>ldlseek(3X)</b>	seek to line number entries of a section of a COFF file
<b>ldnrseek</b>	<b>ldlseek(3X)</b>	seek to line number entries of a section of a COFF file
<b>ldnshread</b>	<b>ldrseek(3X)</b>	seek to relocation entries of a section of a COFF file
<b>ldnsseek</b>	<b>ldshread(3X)</b>	read an indexed/named section header of a COFF file
<b>ldohseek</b>	<b>ldsseek(3X)</b>	seek to an indexed/named section of a COFF file
<b>ldopen</b>	<b>ldohseek(3X)</b>	seek to the optional file header of a COFF file
	<b>ldopen(3X)</b>	open a COFF file for reading

<b>ldrseek</b>	<b>ldrseek(3X)</b>	seek to relocation entries of a section of a COFF file
<b>ldshread</b>	<b>ldshread(3X)</b>	read an indexed/named section header of a COFF file
<b>ldsseek</b>	<b>ldsseek(3X)</b>	seek to an indexed/named section of a COFF file
<b>ldtbindx</b>	<b>ldtbindx(3X)</b>	compute the index of a symbol table entry of a COFF file
<b>ldtbread</b>	<b>ldtbread(3X)</b>	read an indexed symbol table entry of a COFF file
<b>ldtbseek</b>	<b>ldtbseek(3X)</b>	seek to the symbol table of a COFF file
<b>lfind</b>	<b>lsearch(3)</b>	linear search and update
<b>line</b>	<b>plot(3X)</b>	graphics interface
<b>linemod</b>	<b>plot(3X)</b>	graphics interface
<b>localdtconv</b>	<b>localdtconv(3)</b>	get date and time formatting conventions
<b>localeconv</b>	<b>localeconv(3)</b>	get numeric and monetary formatting conventions
<b>localtime</b>	<b>ctime(3V)</b>	convert date and time
<b>lockf</b>	<b>lockf(3)</b>	record locking on files
<b>longjmp</b>	<b>setjmp(3V)</b>	non-local goto
<b>lrand48</b>	<b>drand48(3)</b>	generate uniformly distributed pseudo-random numbers
<b>lsearch</b>	<b>lsearch(3)</b>	linear search and update
<b>lto3</b>	<b>l3tol(3C)</b>	convert between 3-byte integers and long integers
<b>madd</b>	<b>mp(3X)</b>	multiple precision integer arithmetic
<b>madvise</b>	<b>madvise(3)</b>	provide advice to VM system
<b>malloc</b>	<b>malloc(3V)</b>	memory allocator
<b>malloc_debug</b>	<b>malloc(3V)</b>	memory allocator
<b>malloc_verify</b>	<b>malloc(3V)</b>	memory allocator
<b>malloccmap</b>	<b>malloc(3V)</b>	memory allocator
<b>mblen</b>	<b>mblen(3)</b>	multibyte character handling
<b>mbstowcs</b>	<b>mblen(3)</b>	multibyte character handling
<b>mbtowc</b>	<b>mblen(3)</b>	multibyte character handling
<b>ncmp</b>	<b>mp(3X)</b>	multiple precision integer arithmetic
<b>mdiv</b>	<b>mp(3X)</b>	multiple precision integer arithmetic
<b>memalign</b>	<b>malloc(3V)</b>	memory allocator
<b>memccpy</b>	<b>memory(3)</b>	memory operations
<b>memchr</b>	<b>memory(3)</b>	memory operations
<b>memcmp</b>	<b>memory(3)</b>	memory operations
<b>memcpy</b>	<b>memory(3)</b>	memory operations
<b>memory</b>	<b>memory(3)</b>	memory operations
<b>memset</b>	<b>memory(3)</b>	memory operations
<b>mfree</b>	<b>mp(3X)</b>	multiple precision integer arithmetic
<b>min</b>	<b>mp(3X)</b>	multiple precision integer arithmetic
<b>mkstemp</b>	<b>mktemp(3)</b>	make a unique file name
<b>mktemp</b>	<b>mktemp(3)</b>	make a unique file name
<b>mlock</b>	<b>mlock(3)</b>	lock (or unlock) pages in memory
<b>mlockall</b>	<b>mlockall(3)</b>	lock (or unlock) address space
<b>moncontrol</b>	<b>monitor(3)</b>	prepare execution profile
<b>monitor</b>	<b>monitor(3)</b>	prepare execution profile
<b>monstartup</b>	<b>monitor(3)</b>	prepare execution profile
<b>mout</b>	<b>mp(3X)</b>	multiple precision integer arithmetic
<b>move</b>	<b>plot(3X)</b>	graphics interface
<b>mp</b>	<b>mp(3X)</b>	multiple precision integer arithmetic
<b>mrnd48</b>	<b>drand48(3)</b>	generate uniformly distributed pseudo-random numbers
<b>msub</b>	<b>mp(3X)</b>	multiple precision integer arithmetic
<b>msync</b>	<b>msync(3)</b>	synchronize memory with physical storage
<b>mtx</b>	<b>mp(3X)</b>	multiple precision integer arithmetic
<b>mult</b>	<b>mp(3X)</b>	multiple precision integer arithmetic
<b>munlock</b>	<b>mlock(3)</b>	lock (or unlock) pages in memory

<b>munlockall</b>	<b>mlockall(3)</b>	lock (or unlock) address space
<b>ndbm</b>	<b>ndbm(3)</b>	data base subroutines
<b>netname2host</b>	<b>secure_rpc(3N)</b>	library routines for secure remote procedure calls
<b>netname2user</b>	<b>secure_rpc(3N)</b>	library routines for secure remote procedure calls
<b>nextkey</b>	<b>dbm(3X)</b>	data base subroutines
<b>nice</b>	<b>nice(3V)</b>	change nice value of a process
<b>nl_init</b>	<b>setlocale(3V)</b>	set international environment
<b>nl_langinfo</b>	<b>nl_langinfo(3C)</b>	language information
<b>nlist</b>	<b>nlist(3V)</b>	get entries from symbol table
<b>rand48</b>	<b>drand48(3)</b>	generate uniformly distributed pseudo-random numbers
<b>ntohl</b>	<b>byteorder(3N)</b>	convert values between host and network byte order
<b>ntohs</b>	<b>byteorder(3N)</b>	convert values between host and network byte order
<b>on_exit</b>	<b>on_exit(3)</b>	name termination handler
<b>opendir</b>	<b>directory(3V)</b>	directory operations
<b>openlog</b>	<b>syslog(3)</b>	control system log
<b>openpl</b>	<b>plot(3X)</b>	graphics interface
<b>optarg</b>	<b>getopt(3)</b>	get option letter from argument vector
<b>optind</b>	<b>getopt(3)</b>	get option letter from argument vector
<b>passwd2des</b>	<b>xcrypt(3R)</b>	hex encryption and utility routines
<b>pause</b>	<b>pause(3V)</b>	stop until signal
<b>pclose</b>	<b>popen(3S)</b>	open or close a pipe (for I/O) from or to a process
<b>perror</b>	<b>perror(3)</b>	system error messages
<b>plock</b>	<b>plock(3)</b>	lock process, text, or data segment in memory
<b>plot</b>	<b>plot(3X)</b>	graphics interface
<b>point</b>	<b>plot(3X)</b>	graphics interface
<b>popen</b>	<b>popen(3S)</b>	open or close a pipe (for I/O) from or to a process
<b>pow</b>	<b>mp(3X)</b>	multiple precision integer arithmetic
<b>printf</b>	<b>printf(3V)</b>	formatted output conversion
<b>prof</b>	<b>prof(3)</b>	profile within a function
<b>psignal</b>	<b>psignal(3)</b>	system signal messages
<b>publickey</b>	<b>publickey(3R)</b>	get public or secret key
<b>putc</b>	<b>putc(3S)</b>	put character or word on a stream
<b>putchar</b>	<b>putc(3S)</b>	put character or word on a stream
<b>putenv</b>	<b>putenv(3)</b>	change or add value to environment
<b>putpwent</b>	<b>putpwent(3)</b>	write password file entry
<b>puts</b>	<b>puts(3S)</b>	put a string on a stream
<b>putw</b>	<b>putc(3S)</b>	put character or word on a stream
<b>pwdauth</b>	<b>pwdauth(3)</b>	password authentication routines
<b>qsort</b>	<b>qsort(3)</b>	quicker sort
<b>rand</b>	<b>rand(3V)</b>	simple random number generator
<b>random</b>	<b>random(3)</b>	better random number generator
<b>rcmd</b>	<b>rcmd(3N)</b>	routines for returning a stream to a remote command
<b>re_comp</b>	<b>regex(3)</b>	regular expression handler
<b>re_exec</b>	<b>regex(3)</b>	regular expression handler
<b>readdir</b>	<b>directory(3V)</b>	directory operations
<b>realloc</b>	<b>malloc(3V)</b>	memory allocator
<b>realpath</b>	<b>realpath(3)</b>	return the canonicalized absolute pathname
<b>regex</b>	<b>regex(3)</b>	regular expression handler
<b>regexp</b>	<b>regexp(3)</b>	regular expression compile and match routines
<b>registerrpc</b>	<b>rpc_svc_calls(3N)</b>	library routines for registering servers
<b>remexportent</b>	<b>exportent(3)</b>	get exported file system information
<b>remque</b>	<b>insque(3)</b>	insert/remove element from a queue
<b>res_init</b>	<b>resolver(3)</b>	resolver routines

res_mkquery	resolver(3)	resolver routines
res_send	resolver(3)	resolver routines
resolver	resolver(3)	resolver routines
rewind	fseek(3S)	reposition a stream
rewinddir	directory(3V)	directory operations
rexec	rexec(3N)	return stream to a remote command
rindex	string(3)	string operations
rpc	rpc(3N)	library routines for remote procedure calls
rpc_createrr	rpc_clnt_create(3N)	library routines creating and manipulating CLIENT handles
rpow	mp(3X)	multiple precision integer arithmetic
rresvport	rcmd(3N)	routines for returning a stream to a remote command
rtime	rtime(3N)	get remote time
ruserok	rcmd(3N)	routines for returning a stream to a remote command
scandir	scandir(3)	scan a directory
scanf	scanf(3V)	formatted input conversion
seconvert	econvert(3)	output conversion
seed48	drand48(3)	generate uniformly distributed pseudo-random numbers
seekdir	directory(3V)	directory operations
setac	getacinfo(3)	get audit control file information
setbuf	setbuf(3V)	assign buffering to a stream
setbuffer	setbuf(3V)	assign buffering to a stream
setegid	setuid(3V)	set user and group ID
seteuid	setuid(3V)	set user and group ID
setexportent	exportent(3)	get exported file system information
setfsent	getfsent(3)	get file system descriptor file entry
setgid	setuid(3V)	set user and group ID
setgraent	getgraent(3)	get group adjunct file entry
setgrent	getgrent(3V)	get group file entry
sethostent	gethostent(3N)	get network host entry
setjmp	setjmp(3V)	non-local goto
setkey	crypt(3)	password and data encryption
setlinebuf	setbuf(3V)	assign buffering to a stream
setlocale	setlocale(3V)	set international environment
setlogmask	syslog(3)	control system log
setmntent	getmntent(3)	get file system descriptor file entry
setnetent	getnetent(3N)	get network entry
setnetgrent	getnetgrent(3N)	get network group entry
setprotoent	getprotoent(3N)	get protocol entry
setpwaent	getpwaent(3)	get password adjunct file entry
setpwent	getpwent(3V)	get password file entry
setpwfile	getpwent(3V)	get password file entry
setrgid	setuid(3V)	set user and group ID
setrpcent	getrpcent(3N)	get RPC entry
setruid	setuid(3V)	set user and group ID
setservent	getservent(3N)	get service entry
setstate	random(3)	better random number generator
settyent	gettyent(3)	get ttytab file entry
setuid	setuid(3V)	set user and group ID
setusershell	getusershell(3)	get legal user shells
setvbuf	setbuf(3V)	assign buffering to a stream
sfconvert	econvert(3)	output conversion
sgconvert	econvert(3)	output conversion
sigaction	sigaction(3V)	examine and change signal action



<b>sigaddset</b>	<b>sigsetops(3V)</b>	manipulate signal sets
<b>sigdelset</b>	<b>sigsetops(3V)</b>	manipulate signal sets
<b>sigemptyset</b>	<b>sigsetops(3V)</b>	manipulate signal sets
<b>sigfillset</b>	<b>sigsetops(3V)</b>	manipulate signal sets
<b>sigfpe</b>	<b>sigfpe(3)</b>	signal handling for specific SIGFPE codes
<b>siginterrupt</b>	<b>siginterrupt(3V)</b>	allow signals to interrupt system calls
<b>sigismember</b>	<b>sigsetops(3V)</b>	manipulate signal sets
<b>siglongjmp</b>	<b>setjmp(3V)</b>	non-local goto
<b>signal</b>	<b>signal(3V)</b>	simplified software signal facilities
<b>sigsetjmp</b>	<b>setjmp(3V)</b>	non-local goto
<b>sigsetops</b>	<b>sigsetops(3V)</b>	manipulate signal sets
<b>single_to_decimal</b>	<b>floating_to_decimal(3)</b>	convert floating-point value to decimal record
<b>sleep</b>	<b>sleep(3V)</b>	suspend execution for interval
<b>space</b>	<b>plot(3X)</b>	graphics interface
<b>sprintf</b>	<b>printf(3V)</b>	formatted output conversion
<b>srand48</b>	<b>drand48(3)</b>	generate uniformly distributed pseudo-random numbers
<b>srand</b>	<b>rand(3V)</b>	simple random number generator
<b>random</b>	<b>random(3)</b>	better random number generator
<b>scanf</b>	<b>scanf(3V)</b>	formatted input conversion
<b>ssignal</b>	<b>ssignal(3)</b>	software signals
<b>stdio</b>	<b>stdio(3V)</b>	standard buffered input/output package
<b>store</b>	<b>dbm(3X)</b>	data base subroutines
<b>strcasecmp</b>	<b>string(3)</b>	string operations
<b>strcat</b>	<b>string(3)</b>	string operations
<b>strchr</b>	<b>string(3)</b>	string operations
<b>strcmp</b>	<b>string(3)</b>	string operations
<b>strcoll</b>	<b>strcoll(3)</b>	compare or transform strings using collating information
<b>strcpy</b>	<b>string(3)</b>	string operations
<b>strcspn</b>	<b>string(3)</b>	string operations
<b>strdup</b>	<b>string(3)</b>	string operations
<b>strftime</b>	<b>ctime(3V)</b>	convert date and time
<b>string_to_decimal</b>	<b>string_to_decimal(3)</b>	parse characters into decimal record
<b>strlen</b>	<b>string(3)</b>	string operations
<b>strncasecmp</b>	<b>string(3)</b>	string operations
<b>strncat</b>	<b>string(3)</b>	string operations
<b>strncmp</b>	<b>string(3)</b>	string operations
<b>strncpy</b>	<b>string(3)</b>	string operations
<b>strpbrk</b>	<b>string(3)</b>	string operations
<b>strptime</b>	<b>ctime(3V)</b>	convert date and time
<b>strrchr</b>	<b>string(3)</b>	string operations
<b>strspn</b>	<b>string(3)</b>	string operations
<b>strstr</b>	<b>string(3)</b>	string operations
<b>strtod</b>	<b>strtod(3)</b>	convert string to double-precision number
<b>strtok</b>	<b>string(3)</b>	string operations
<b>strtol</b>	<b>strtol(3)</b>	convert string to integer
<b>strxfrm</b>	<b>strcoll(3)</b>	compare or transform strings using collating information
<b>stty</b>	<b>stty(3C)</b>	set and get terminal state
<b>svc_destroy</b>	<b>rpc_svc_create(3N)</b>	library routines for dealing with the creation of server handles
<b>svc_fds</b>	<b>rpc_svc_reg(3N)</b>	library routines for RPC servers
<b>svc_fdset</b>	<b>rpc_svc_reg(3N)</b>	library routines for RPC servers
<b>svc_freeargs</b>	<b>rpc_svc_reg(3N)</b>	library routines for RPC servers
<b>svc_getargs</b>	<b>rpc_svc_reg(3N)</b>	library routines for RPC servers
<b>svc_getcaller</b>	<b>rpc_svc_reg(3N)</b>	library routines for RPC servers

<b>svc_getreq</b>	<b>rpc_svc_reg(3N)</b>	library routines for RPC servers
<b>svc_getreqset</b>	<b>rpc_svc_reg(3N)</b>	library routines for RPC servers
<b>svc_register</b>	<b>rpc_svc_calls(3N)</b>	library routines for registering servers
<b>svc_run</b>	<b>rpc_svc_reg(3N)</b>	library routines for RPC servers
<b>svc_sendreply</b>	<b>rpc_svc_reg(3N)</b>	library routines for RPC servers
<b>svc_unregister</b>	<b>rpc_svc_calls(3N)</b>	library routines for registering servers
<b>svcerr_auth</b>	<b>rpc_svc_err(3N)</b>	library routines for server side remote procedure call errors
<b>svcerr_decode</b>	<b>rpc_svc_err(3N)</b>	library routines for server side remote procedure call errors
<b>svcerr_noproc</b>	<b>rpc_svc_err(3N)</b>	library routines for server side remote procedure call errors
<b>svcerr_noprogram</b>	<b>rpc_svc_err(3N)</b>	library routines for server side remote procedure call errors
<b>svcerr_progvers</b>	<b>rpc_svc_err(3N)</b>	library routines for server side remote procedure call errors
<b>svcerr_systemerr</b>	<b>rpc_svc_err(3N)</b>	library routines for server side remote procedure call errors
<b>svcerr_weakauth</b>	<b>rpc_svc_err(3N)</b>	library routines for server side remote procedure call errors
<b>svcfld_create</b>	<b>rpc_svc_create(3N)</b>	library routines for dealing with the creation of server handles
<b>svcrow_create</b>	<b>rpc_svc_create(3N)</b>	library routines for dealing with the creation of server handles
<b>svctcp_create</b>	<b>rpc_svc_create(3N)</b>	library routines for dealing with the creation of server handles
<b>svcudp_bufcreate</b>	<b>rpc_svc_create(3N)</b>	library routines for dealing with the creation of server handles
<b>swab</b>	<b>swab(3)</b>	swap bytes
<b>sys_siglist</b>	<b>psignal(3)</b>	system signal messages
<b>syslog</b>	<b>syslog(3)</b>	control system log
<b>system</b>	<b>system(3)</b>	issue a shell command
<b>t_accept</b>	<b>t_accept(3N)</b>	accept a connect request
<b>t_alloc</b>	<b>t_alloc(3N)</b>	allocate a library structure
<b>t_bind</b>	<b>t_bind(3N)</b>	bind an address to a transport endpoint
<b>t_close</b>	<b>t_close(3N)</b>	close a transport endpoint
<b>t_connect</b>	<b>t_connect(3N)</b>	establish a connection with another transport user
<b>t_error</b>	<b>t_error(3N)</b>	produce error message
<b>t_free</b>	<b>t_free(3N)</b>	free a library structure
<b>t_getinfo</b>	<b>t_getinfo(3N)</b>	get protocol-specific service information
<b>t_getstate</b>	<b>t_getstate(3N)</b>	get the current state
<b>t_listen</b>	<b>t_listen(3N)</b>	listen for a connect request
<b>t_look</b>	<b>t_look(3N)</b>	look at the current event on a transport endpoint
<b>t_open</b>	<b>t_open(3N)</b>	establish a transport endpoint
<b>t_optmgmt</b>	<b>t_optmgmt(3N)</b>	manage options for a transport endpoint
<b>t_rcv</b>	<b>t_rcv(3N)</b>	receive normal or expedited data sent over a connection
<b>t_rcvconnect</b>	<b>t_rcvconnect(3N)</b>	receive the confirmation from a connect request
<b>t_rcvdis</b>	<b>t_rcvdis(3N)</b>	retrieve information from disconnect
<b>t_rcvrel</b>	<b>t_rcvrel(3N)</b>	acknowledge receipt of an orderly release indication
<b>t_rcvudata</b>	<b>t_rcvudata(3N)</b>	receive a data unit
<b>t_rcvuderr</b>	<b>t_rcvuderr(3N)</b>	receive a unit data error indication
<b>t_snd</b>	<b>t_snd(3N)</b>	send normal or expedited data over a connection
<b>t_snddis</b>	<b>t_snddis(3N)</b>	send user-initiated disconnect request
<b>t_sndrel</b>	<b>t_sndrel(3N)</b>	initiate an orderly release
<b>t_sndudata</b>	<b>t_sndudata(3N)</b>	send a data unit
<b>t_sync</b>	<b>t_sync(3N)</b>	synchronize transport library
<b>t_unbind</b>	<b>t_unbind(3N)</b>	disable a transport endpoint
<b>tcdrain</b>	<b>termios(3V)</b>	terminal control functions
<b>tcflow</b>	<b>termios(3V)</b>	terminal control functions
<b>tcflush</b>	<b>termios(3V)</b>	terminal control functions
<b>tcgetattr</b>	<b>termios(3V)</b>	terminal control functions
<b>tcgetpgrp</b>	<b>tcgetpgrp(3V)</b>	get, set foreground process group ID
<b>tcsendbreak</b>	<b>termios(3V)</b>	terminal control functions
<b>tcsetattr</b>	<b>termios(3V)</b>	terminal control functions

<b>tcsetpgrp</b>	<b>tcgetpgrp(3V)</b>	get, set foreground process group ID
<b>tdelete</b>	<b>tsearch(3)</b>	manage binary search trees
<b>telldir</b>	<b>directory(3V)</b>	directory operations
<b>tempnam</b>	<b>tmpnam(3S)</b>	create a name for a temporary file
<b>termcap</b>	<b>termcap(3X)</b>	terminal independent operation routines
<b>termios</b>	<b>termios(3V)</b>	terminal control functions
<b>textdomain</b>	<b>gettext(3)</b>	retrieve a message string, get and set text domain
<b>tfind</b>	<b>tsearch(3)</b>	manage binary search trees
<b>tgetent</b>	<b>termcap(3X)</b>	terminal independent operation routines
<b>tgetflag</b>	<b>termcap(3X)</b>	terminal independent operation routines
<b>tgetnum</b>	<b>termcap(3X)</b>	terminal independent operation routines
<b>tgetstr</b>	<b>termcap(3X)</b>	terminal independent operation routines
<b>tgoto</b>	<b>termcap(3X)</b>	terminal independent operation routines
<b>time</b>	<b>time(3V)</b>	get date and time
<b>timegm</b>	<b>ctime(3V)</b>	convert date and time
<b>timelocal</b>	<b>ctime(3V)</b>	convert date and time
<b>times</b>	<b>times(3V)</b>	get process times
<b>timezone</b>	<b>timezone(3C)</b>	get time zone name given offset from GMT
<b>tmpfile</b>	<b>tmpfile(3S)</b>	create a temporary file
<b>tmpnam</b>	<b>tmpnam(3S)</b>	create a name for a temporary file
<b>toascii</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>tolower</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>toupper</b>	<b>ctype(3V)</b>	character classification and conversion macros and functions
<b>tputs</b>	<b>termcap(3X)</b>	terminal independent operation routines
<b>tsearch</b>	<b>tsearch(3)</b>	manage binary search trees
<b>ttyname</b>	<b>ttyname(3V)</b>	find name of a terminal
<b>ttyslot</b>	<b>ttyslot(3V)</b>	find the slot in the utmp file of the current process
<b>twalk</b>	<b>tsearch(3)</b>	manage binary search trees
<b>tzset</b>	<b>ctime(3V)</b>	convert date and time
<b>tzsetwall</b>	<b>ctime(3V)</b>	convert date and time
<b>ualarm</b>	<b>ualarm(3)</b>	schedule signal after interval in microseconds
<b>ulimit</b>	<b>ulimit(3C)</b>	get and set user limits
<b>ungetc</b>	<b>ungetc(3S)</b>	push character back into input stream
<b>user2netname</b>	<b>secure_rpc(3N)</b>	library routines for secure remote procedure calls
<b>usleep</b>	<b>usleep(3)</b>	suspend execution for interval in microseconds
<b>utime</b>	<b>utime(3V)</b>	set file times
<b>valloc</b>	<b>malloc(3V)</b>	memory allocator
<b>values</b>	<b>values(3)</b>	machine-dependent values
<b>varargs</b>	<b>varargs(3)</b>	handle variable argument list
<b>vfprintf</b>	<b>vprintf(3V)</b>	print formatted output of a varargs argument list
<b>vlimit</b>	<b>vlimit(3C)</b>	control maximum system resource consumption
<b>vprintf</b>	<b>vprintf(3V)</b>	print formatted output of a varargs argument list
<b>vsprintf</b>	<b>vprintf(3V)</b>	print formatted output of a varargs argument list
<b>vsyslog</b>	<b>vsyslog(3)</b>	log message with a varargs argument list
<b>vtimes</b>	<b>vtimes(3C)</b>	get information about resource utilization
<b>wcstombs</b>	<b>mblen(3)</b>	multibyte character handling
<b>wctomb</b>	<b>mblen(3)</b>	multibyte character handling
<b>xcrypt</b>	<b>xcrypt(3R)</b>	hex encryption and utility routines
<b>xdecrypt</b>	<b>xcrypt(3R)</b>	hex encryption and utility routines
<b>xdr</b>	<b>xdr(3N)</b>	library routines for external data representation
<b>xdr_accepted_reply</b>	<b>rpc_xdr(3N)</b>	XDR library routines for remote procedure calls
<b>xdr_array</b>	<b>xdr_complex(3N)</b>	library routines for translating complex data types
<b>xdr_authunix_parms</b>	<b>rpc_xdr(3N)</b>	XDR library routines for remote procedure calls

xdr_bool	xdr_simple(3N)	library routines for translating simple data types
xdr_bytes	xdr_complex(3N)	library routines for translating complex data types
xdr_callhdr	rpc_xdr(3N)	XDR library routines for remote procedure calls
xdr_callmsg	rpc_xdr(3N)	XDR library routines for remote procedure calls
xdr_char	xdr_simple(3N)	library routines for translating simple data types
xdr_destroy	xdr_create(3N)	library routines for XDR stream creation
xdr_double	xdr_simple(3N)	library routines for translating simple data types
xdr_enum	xdr_simple(3N)	library routines for translating simple data types
xdr_float	xdr_simple(3N)	library routines for translating simple data types
xdr_free	xdr_simple(3N)	library routines for translating simple data types
xdr_getpos	xdr_admin(3N)	library routines for management of the XDR stream
xdr_inline	xdr_admin(3N)	library routines for management of the XDR stream
xdr_int	xdr_simple(3N)	library routines for translating simple data types
xdr_long	xdr_simple(3N)	library routines for translating simple data types
xdr_opaque	xdr_complex(3N)	library routines for translating complex data types
xdr_opaque_auth	rpc_xdr(3N)	XDR library routines for remote procedure calls
xdr_pamp	portmap(3N)	library routines for RPC bind service
xdr_pmaplist	portmap(3N)	library routines for RPC bind service
xdr_pointer	xdr_complex(3N)	library routines for translating complex data types
xdr_reference	xdr_complex(3N)	library routines for translating complex data types
xdr_rejected_reply	rpc_xdr(3N)	XDR library routines for remote procedure calls
xdr_replymsg	rpc_xdr(3N)	XDR library routines for remote procedure calls
xdr_setpos	xdr_admin(3N)	library routines for management of the XDR stream
xdr_short	xdr_simple(3N)	library routines for translating simple data types
xdr_string	xdr_complex(3N)	library routines for translating complex data types
xdr_u_char	xdr_simple(3N)	library routines for translating simple data types
xdr_u_int	xdr_simple(3N)	library routines for translating simple data types
xdr_u_long	xdr_simple(3N)	library routines for translating simple data types
xdr_u_short	xdr_simple(3N)	library routines for translating simple data types
xdr_union	xdr_complex(3N)	library routines for translating complex data types
xdr_vector	xdr_complex(3N)	library routines for translating complex data types
xdr_void	xdr_simple(3N)	library routines for translating simple data types
xdr_wrapstring	xdr_complex(3N)	library routines for translating complex data types
xdrmem_create	xdr_create(3N)	library routines for XDR stream creation
xdrrec_create	xdr_create(3N)	library routines for XDR stream creation
xdrrec_endofrecord	xdr_admin(3N)	library routines for management of the XDR stream
xdrrec_eof	xdr_admin(3N)	library routines for management of the XDR stream
xdrrec_readbytes	xdr_admin(3N)	library routines for management of the XDR stream
xdrrec_skiprecord	xdr_admin(3N)	library routines for management of the XDR stream
xdrstdio_create	xdr_create(3N)	library routines for XDR stream creation
xencrypt	xcrypt(3R)	hex encryption and utility routines
xprt_register	rpc_svc_calls(3N)	library routines for registering servers
xprt_unregister	rpc_svc_calls(3N)	library routines for registering servers
xtom	mp(3X)	multiple precision integer arithmetic
yp_all	ypclnt(3N)	NIS client interface
yp_bind	ypclnt(3N)	NIS client interface
yp_first	ypclnt(3N)	NIS client interface
yp_get_default_domain	ypclnt(3N)	NIS client interface
yp_master	ypclnt(3N)	NIS client interface
yp_match	ypclnt(3N)	NIS client interface
yp_next	ypclnt(3N)	NIS client interface
yp_order	ypclnt(3N)	NIS client interface
yp_unbind	ypclnt(3N)	NIS client interface

<b>yp_update</b>	<b>ypupdate(3N)</b>	changes NIS information
<b>ypclnt</b>	<b>ypclnt(3N)</b>	NIS client interface
<b>yperr_string</b>	<b>ypclnt(3N)</b>	NIS client interface
<b>ypprot_err</b>	<b>ypclnt(3N)</b>	NIS client interface

**NAME**

a64l, l64a – convert between long integer and base-64 ASCII string

**SYNOPSIS**

**long a64l(s)**

**char \*s;**

**char \*l64a(l)**

**long l;**

**DESCRIPTION**

These functions are used to maintain numbers stored in *base-64* ASCII characters. This is a notation by which long integers can be represented by up to six characters; each character represents a “digit” in a radix-64 notation.

The characters used to represent “digits” are ‘.’ for 0, ‘/’ for 1, 0 through 9 for 2–11, A through Z for 12–37, and a through z for 38–63.

**a64l()** takes a pointer to a null-terminated base-64 representation and returns a corresponding **long** value. If the string pointed to by *s* contains more than six characters, **a64l()** will use the first six.

**l64a()** takes a **long** argument and returns a pointer to the corresponding base-64 representation. If the argument is 0, **l64a()** returns a pointer to a null string.

**BUGS**

The value returned by **l64a()** is a pointer into a static buffer, the contents of which are overwritten by each call.

**NAME**

**abort** – generate a fault

**SYNOPSIS**

**abort()**

**DESCRIPTION**

**abort()** first closes all open files if possible, then sends an IOT signal to the process. This signal usually results in termination with a core dump, which may be used for debugging.

It is possible for **abort()** to return control if SIGIOT is caught or ignored, in which case the value returned is that of the **kill(2V)** system call.

**SEE ALSO**

**adb(1)**, **exit(2V)**, **kill(2V)**, **signal(3V)**

**DIAGNOSTICS**

If SIGIOT is neither caught nor ignored, and the current directory is writable, a core dump is produced and the message '**abort – core dumped**' is written by the shell.

**NAME**

**abs** – integer absolute value

**SYNOPSIS**

```
abs(i)  
int i;
```

**DESCRIPTION**

**abs()** returns the absolute value of its integer operand.

**SEE ALSO**

**ieee\_functions(3M)** for **fabs()**

**BUGS**

Applying the **abs()** function to the most negative integer generates a result which is the most negative integer. That is, **abs(0x80000000)** returns **0x80000000** as a result.



**NAME**

**aio\_cancel** – cancel an asynchronous operation

**SYNOPSIS**

```
#include <sys/asynch.h>
```

```
int aio_cancel(resultp)
```

```
    aio_result_t *resultp;
```

**DESCRIPTION**

**aio\_cancel()** cancels the asynchronous operation associated with the result buffer pointed to by *resultp*. It may not be possible to immediately cancel an operation which is in progress and in this case, **aio\_cancel()** will not wait to cancel it.

Upon successful completion, **aio\_cancel()** will return 0 and the requested operation will be canceled. The application will not receive the **SIGIO** completion signal for an asynchronous operation which is successfully canceled.

**RETURN VALUES**

**aio\_cancel()** returns:

0        on success.

-1       on failure and sets **errno** to indicate the error.

**ERRORS**

**aio\_cancel()** will fail if any of the following are true:

**EACCES**        The parameter *resultp* does not correspond to an outstanding asynchronous operation.  
                 The operation could not be cancelled.

**EFAULT**        The parameter *resultp* points to an address that is outside of the address space of the requesting process.

**SEE ALSO**

**aio\_read(3)**, **aiowait(3)**

## NAME

`aioread`, `aiowrite` – asynchronous I/O operations

## SYNOPSIS

```
#include <sys/asynch.h>

int aioread(fd, bufp, bufs, offset, whence, resultp)
int fd;
char *bufp;
int bufs;
int offset;
int whence;
aio_result_t *resultp;

int aiowrite(fd, bufp, bufs, offset, whence, resultp)
int fd;
char *bufp;
int bufs;
int offset;
int whence;
aio_result_t *resultp;
```

## DESCRIPTION

`aioread()` initiates one asynchronous `read(2V)` and returns control to the calling program. The `read()` continues concurrently with other activity of the process. An attempt is made to read *bufs* bytes of data from the object referenced by the descriptor *fd* into the buffer pointed to by *bufp*.

`aiowrite()` initiates one asynchronous `write(2V)` and returns control to the calling program. The `write()` continues concurrently with other activity of the process. An attempt is made to write *bufs* bytes of data from the buffer pointed to by *bufp* to the object referenced by the descriptor *fd*.

On objects capable of seeking, the I/O operation starts at the position specified by *whence* and *offset*. These parameters have the same meaning as the corresponding parameters to the `lseek(2V)` function. On objects not capable of seeking the I/O operation always start from the current position and the parameters *whence* and *offset* are ignored. The seek pointer for objects capable of seeking is not updated by `aioread()` or `aiowrite()`. Sequential asynchronous operations on these devices must be managed by the application using the *whence* and *offset* parameters.

The result of the asynchronous operation is stored in the structure pointed to by *resultp*:

```
int aio_return;      /* return value of read() or write() */
int aio_errno;      /* value of errno for read() or write() */
```

Upon completion of the operation both *aio\_return* and *aio\_errno* are set to reflect the result of the operation. `AIO_INPROGRESS` is not a value used by the system so the client may detect a change in state by initializing *aio\_return* to this value.

Notification of the completion of an asynchronous I/O operation may be obtained synchronously through the `aiowait(3)` function, or asynchronously through the signal mechanism. Asynchronous notification is accomplished by generating the `SIGIO` signal. The delivery of this instance of the `SIGIO` signal is reliable in that a signal delivered while the handler is executing is not lost. If the client ensures that `aiowait(3)` returns nothing (using a polling timeout) before returning from the signal handler, no asynchronous I/O notifications are lost. The `aiowait(3)` function is the only way to dequeue an asynchronous notification. Note: `SIGIO` may have several meanings simultaneously: for example, that a descriptor generated `SIGIO` and an asynchronous operation completed. Further, issuing an asynchronous request successfully guarantees that space exists to queue the completion notification.

`close(2V)`, `exit(2V)` and `execve(2V)` will block until all pending asynchronous I/O operations can be cancelled by the system.

It is an error to use the same result buffer in more than one outstanding request. These structures may only be reused after the system has completed the operation.

**RETURN VALUES**

**aioread()** and **aiowrite()** return:

- 0 on success.
- 1 on failure and set **errno** to indicate the error.

**ERRORS**

- EBADF** *fd* is not a valid file descriptor open for reading.
- EFAULT** At least one of *bufp* or *resultp* points to an address outside the address space of the requesting process.
- EINVAL** The parameter *resultp* is currently being used by an outstanding asynchronous request.
- EPROCLIM** The number of asynchronous requests that the system can handle at any one time has been exceeded

**SEE ALSO**

**close(2V)**, **execve(2V)**, **exit(2V)**, **lseek(2V)**, **open(2V)**, **read(2V)**, **sigvec(2)**, **write(2V)**, **aiocancel(3)**, **aiowait(3)**

**NAME**

**aiowait** – wait for completion of asynchronous I/O operation

**SYNOPSIS**

```
#include <sys/asynch.h>
#include <sys/time.h>

aio_result_t *aiowait(timeout)
struct timeval *timeout;
```

**DESCRIPTION**

**aiowait()** suspends the calling process until one of its outstanding asynchronous I/O operations completes. This provides a synchronous method of notification.

If *timeout* is a non-zero pointer, it specifies a maximum interval to wait for the completion of an asynchronous I/O operation. If *timeout* is a zero pointer, then **aiowait()** blocks indefinitely. To effect a poll, the *timeout* parameter should be non-zero, pointing to a zero-valued *timeval* structure. The *timeval* structure is defined in `<sys/time.h>` as:

```
struct timeval {
    long tv_sec;           /* seconds */
    long tv_usec;        /* and microseconds */
};
```

**NOTES**

**aiowait()** is the only way to dequeue an asynchronous notification. It may be used either inside a SIGIO signal handler or in the main program. Note: one SIGIO signal may represent several queued events.

**RETURN VALUES**

On success, **aiowait()** returns a pointer to the result structure used when the completed asynchronous I/O operation was requested. On failure, it returns `-1` and sets `errno` to indicate the error. **aiowait()** returns `0` if the time limit expires.

**ERRORS**

<b>EFAULT</b>	<i>timeout</i> points to an address outside the address space of the requesting process.
<b>EINTR</b>	A signal was delivered before an asynchronous I/O operation completed. The time limit expired.
<b>EINVAL</b>	There are no outstanding asynchronous I/O requests.

**SEE ALSO**

**aiocancel(3)**, **aioread(3)**

**NAME**

alarm – schedule signal after specified time

**SYNOPSIS**

```
unsigned int alarm(seconds)
unsigned int seconds;
```

**DESCRIPTION**

**alarm()** sends the signal **SIGALRM** (see **sigvec(2)**), to the invoking process after *seconds* seconds. Unless caught or ignored, the signal terminates the process.

**alarm()** requests are not stacked; successive calls reset the alarm clock. If the argument is 0, any **alarm()** request is canceled. Because of scheduling delays, resumption of execution of when the signal is caught may be delayed an arbitrary amount. The longest specifiable delay time is 2147483647 seconds.

The return value is the amount of time previously remaining in the alarm clock.

**SEE ALSO**

**sigpause(2V)**, **sigvec(2)**, **signal(3V)**, **sleep(3V)**, **ualarm(3)**, **usleep(3)**

**WARNINGS**

**alarm()** is slightly incompatible with the default version of **sleep(3V)**. The alarm signal is not sent when one would expect for programs that wait one second of clock time between successive calls to **sleep()**. Each **sleep()** call postpones the alarm signal that would have been sent during the requested sleep period for one second. Use System V **sleep(3V)** to avoid this delay.

**NAME**

assert – program verification

**SYNOPSIS**

```
#include <assert.h>
```

```
assert(expression)
```

**DESCRIPTION**

`assert()` is a macro that indicates *expression* is expected to be true at this point in the program. If *expression* is false (0), it displays a diagnostic message on the standard output and exits (see `exit(2V)`). Compiling with the `cc(1V)` option `-DNDEBUG`, or placing the preprocessor control statement

```
#define NDEBUG
```

before the “`#include <assert.h>`” statement effectively deletes `assert()` from the program.

**SYSTEM V DESCRIPTION**

The System V version of `assert()` calls `abort(3)` rather than `exit()`.

**SEE ALSO**

`cc(1V)`, `exit(2V)`, `abort(3)`

**DIAGNOSTICS**

**Assertion failed: file *f* line *n***

The expression passed to the `assert()` statement at line *n* of source file *f* was false.

**SYSTEM V DIAGNOSTICS**

**Assertion failed: *expression*, file *f*, line *n***

The *expression* passed to the `assert()` statement at line *n* of source file *f* was false.

**NAME**

**audit\_args, audit\_text** – produce text audit message

**SYNOPSIS**

```
#include <sys/label.h>
```

```
#include <sys/audit.h>
```

```
audit_args(event, argc, argv)
```

```
int event;
```

```
int argc;
```

```
char **argv;
```

```
audit_text(event, error, retval, argc, argv)
```

```
int event;
```

```
int error;
```

```
int retval;
```

```
int argc;
```

```
char **argv;
```

**DESCRIPTION**

These functions provide text interfaces to the **audit(2)** system call. In both calls, the *event* parameter identifies the event class of the action, and *argc* is the number of strings found in the vector *argv*. The *error* parameter is used to determine the failure or success of the audited operation. A negative value is always audited. A zero value is audited as a successful event. A positive value is audited as an event failure. The *retval* parameter is the return value or exit code that the invoking program will have.

**audit\_args()** is equivalent to **audit\_text()** with *error* and *retval* parameters of  $-1$ .

**SEE ALSO**

**audit(2)**

**NAME**

**bindresvport** – bind a socket to a privileged IP port

**SYNOPSIS**

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

int bindresvport(sd, sin)
int sd;
struct sockaddr_in *sin;
```

**DESCRIPTION**

**bindresvport()** is used to bind a socket descriptor to a privileged IP port, that is, a port number in the range 0-1023. The routine returns 0 if it is successful, otherwise -1 is returned and **errno** set to reflect the cause of the error. This routine differs with **rresvport** (see **rcmd(3N)**) in that this works for any IP socket, whereas **rresvport()** only works for TCP.

Only root can bind to a privileged port; this call will fail for any other users.

**SEE ALSO**

**rcmd(3N)**



**NAME**

`bsearch` – binary search a sorted table

**SYNOPSIS**

```
#include <search.h>
```

```
char *bsearch ((char *) key, (char *) base, nel, sizeof (*key), compar)  
unsigned nel;  
int (*compar)();
```

**DESCRIPTION**

`bsearch()` is a binary search routine generalized from Knuth (6.2.1) Algorithm B. It returns a pointer into a table indicating where a datum may be found. The table must be previously sorted in increasing order according to a provided comparison function. *key* points to a datum instance to be sought in the table. *base* points to the element at the base of the table. *nel* is the number of elements in the table. *compar* is the name of the comparison function, which is called with two arguments that point to the elements being compared. The function must return an integer less than, equal to, or greater than zero as accordingly the first argument is to be considered less than, equal to, or greater than the second.

**EXAMPLE**

The example below searches a table containing pointers to nodes consisting of a string and its length. The table is ordered alphabetically on the string in the node pointed to by each entry.

This code fragment reads in strings and either finds the corresponding node, in which case it prints out the string and its length, or it prints an error message.

```

#include <stdio.h>
#include <search.h>
#define TABSIZE      1000
struct node {
    char *string;
    int length;
};
struct node table[TABSIZE]; /* table to be searched */
.
.
.
{
    struct node *node_ptr, node;
    int node_compare(); /* routine to compare 2 nodes */
    char str_space[20]; /* space to read string into */
    .
    .
    .
    node.string = str_space;
    while (scanf("%s", node.string) != EOF) {
        node_ptr = (struct node *)bsearch((char *)&node,
            (char *)table, TABSIZE,
            sizeof(struct node), node_compare);
        if (node_ptr != NULL) {
            (void)printf("string = %20s, length = %d\n",
                node_ptr->string, node_ptr->length);
        } else {
            (void)printf("not found: %s\n", node.string);
        }
    }
}
/*
    This routine compares two nodes based on an
    alphabetical ordering of the string field.
*/
int
node_compare(node1, node2)
struct node *node1, *node2;
{
    return strcmp(node1->string, node2->string);
}

```

**NOTES**

The pointers to the key and the element at the base of the table should be of type pointer-to-element, and cast to type pointer-to-character.

The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

Although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

**SEE ALSO**

**hsearch(3), lsearch(3), qsort(3), tsearch(3)**

**DIAGNOSTICS**

A NULL pointer is returned if the key cannot be found in the table.

**NAME**

bstring, bcopy, bcmp, bzero, ffs – bit and byte string operations

**SYNOPSIS**

```
void
bcopy(b1, b2, length)
char *b1, *b2;
int length;

int bcmp(b1, b2, length)
char *b1, *b2;
int length;

void
bzero(b, length)
char *b;
int length;

int ffs(i)
int i;
```

**DESCRIPTION**

The functions **bcopy**, **bcmp**, and **bzero()** operate on variable length strings of bytes. They do not check for null bytes as the routines in **string(3)** do.

**bcopy()** copies *length* bytes from string *b1* to the string *b2*. Overlapping strings are handled correctly.

**bcmp()** compares byte string *b1* against byte string *b2*, returning zero if they are identical, non-zero otherwise. Both strings are assumed to be *length* bytes long. **bcmp()** of length zero bytes always returns zero.

**bzero()** places *length* 0 bytes in the string *b*.

**ffs()** finds the first bit set in the argument passed it and returns the index of that bit. Bits are numbered starting at 1 from the right. A return value of zero indicates that the value passed is zero.

**NOTES**

The **bcmp()** and **bcopy()** routines take parameters backwards from **strcmp()** and **strcpy()**.

**SEE ALSO**

**string(3)**

**NAME**

byteorder, htonl, htons, ntohl, ntohs – convert values between host and network byte order

**SYNOPSIS**

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

netlong = htonl(hostlong);
u_long netlong, hostlong;

netshort = htons(hostshort);
u_short netshort, hostshort;

hostlong = ntohl(netlong);
u_long hostlong, netlong;

hostshort = ntohs(netshort);
u_short hostshort, netshort;
```

**DESCRIPTION**

These routines convert 16 and 32 bit quantities between network byte order and host byte order. On Sun-2, Sun-3 and Sun-4 systems, these routines are defined as NULL macros in the include file `<netinet/in.h>`. On Sun386i systems, these routines are functional since its host byte order is different from network byte order.

These routines are most often used in conjunction with Internet addresses and ports as returned by `gethostent(3N)` and `getservent(3N)`.

**SEE ALSO**

`gethostent(3N)`, `getservent(3N)`

**NAME**

**catgets, catgetmsg** – get message from a message catalog

**SYNOPSIS**

```
#include <nl_types.h>
```

```
char *catgets(catd, set_num, msg_num, s)
```

```
nl_catd catd;
```

```
int set_num, msg_num;
```

```
char *s;
```

```
char *catgetmsg(catd, set_num, msg_num, buf, buflen)
```

```
nl_catd catd;
```

```
int set_num;
```

```
int msg_num;
```

```
int buflen;
```

**DESCRIPTION**

**catgets()** reads the message *msg\_num*, in set *set\_num*, from the message catalog identified by *catd*. *catd* is a catalog descriptor returned from an earlier call to **catopen(3C)**. *s* points to a default message string which will be returned by **catgets()** if the identified message catalog is not currently available. The message-text is contained in an internal buffer area and should be copied by the application if it is to be saved or re-used after further calls to **catgets()**.

**catgetmsg()** attempts to read up to *buflen* - 1 bytes of a message string into the area pointed to by *buf*. *buflen* is an integer value containing the size in bytes of *buf*. The return string is always terminated with a null byte.

**RETURN VALUES**

On success, **catgets()** returns a pointer to an internal buffer area containing the null-terminated message string. **catgets()** returns a pointer to *s* if it fails because the message catalog specified by *catd* is not currently available. Otherwise, **catgets()** returns a pointer to an empty string if the message catalog is available but does not contain the specified message.

On success, **catgetmsg()** returns a pointer to the message string in *buf*. If *catd* is invalid or if *set\_num* or *msg\_num* is not in the message catalog, **catgetmsg()** returns a pointer to an empty string.

**SEE ALSO**

**catopen(3C), locale(5)**

**NAME**

catopen, catclose – open/close a message catalog

**SYNOPSIS**

```
#include <nl_types.h>

nl_catd catopen(name, oflag)
char *name;
int oflag;

int catclose(catd)
nl_catd catd;
```

**DESCRIPTION**

**catopen()** opens a message catalog and returns a catalog descriptor. *name* specifies the name of the message catalog to be opened. If *name* contains a '/' then *name* specifies a pathname for the message catalog. Otherwise, the environment variable NLSPATH is used with *name* substituted for %N (see **locale(5)**). If NLSPATH does not exist in the environment, or if a message catalog cannot be opened in any of the paths specified by NLSPATH, the `/etc/locale/LC_MESSAGES/locale` directory is searched for a message catalog with filename *name*, followed by the `/usr/share/lib/locale/LC_MESSAGES/locale` directory. In both cases *locale* stands for the current setting of the LC\_MESSAGES category of locale.

*oflag* is reserved for future use and should be set to 0 (zero). The results of setting this field to any other value are undefined.

**catclose()** closes the message catalog identified by *catd*. It invalidates any following references to the message catalog defined by *catd*.

**RETURN VALUES**

**catopen()** returns a message catalog descriptor on success. On failure, it returns -1.

**catclose()** returns:

0        on success.  
-1       on failure.

**SEE ALSO**

**catgets(3C)**, **locale(5)**

**NOTES**

Using **catopen()** and **catclose()** in conjunction with **gettext()** or **textdomain()** (see **gettext(3)**) is undefined.

**NAME**

clock – report CPU time used

**SYNOPSIS**

**long** clock ( )

**DESCRIPTION**

**clock()** returns the amount of CPU time (in microseconds) used since the first call to **clock**. The time reported is the sum of the user and system times of the calling process and its terminated child processes for which it has executed **wait(2V)** or **system(3)**.

The resolution of the clock is 16.667 milliseconds.

**SEE ALSO**

**wait(2V)**, **system(3)**, **times(3V)**

**BUGS**

The value returned by **clock()** is defined in microseconds for compatibility with systems that have CPU clocks with much higher resolution. Because of this, the value returned will wrap around after accumulating only 2147 seconds of CPU time (about 36 minutes).



## NAME

`crypt`, `_crypt`, `setkey`, `encrypt` – password and data encryption

## SYNOPSIS

```
char *crypt(key, salt)
char *key, *salt;

char *_crypt(key, salt)
char *key, *salt;

setkey(key)
char *key;

encrypt(block, edflag)
char *block;
```

## DESCRIPTION

`crypt()` is the password encryption routine, based on the NBS Data Encryption Standard, with variations intended (among other things) to frustrate use of hardware implementations of the DES for key search.

The first argument to `crypt()` is normally a user's typed password. The second is a 2-character string chosen from the set `[a-zA-Z0-9./]`. Unless it starts with `'##'` or `'#$'`, the *salt* string is used to perturb the DES algorithm in one of 4096 different ways, after which the password is used as the key to encrypt repeatedly a constant string. The returned value points to the encrypted password, in the same alphabet as the salt. The first two characters are the salt itself.

If the *salt* string starts with `'##'`, `pwdauth(3)` is called. If `pwdauth` returns TRUE, the salt is returned from `crypt`. Otherwise, NULL is returned. If the *salt* string starts with `'#$'`, `grpauth` (see `pwdauth(3)`) is called. If `grpauth` returns TRUE, the salt is returned from `crypt`. Otherwise, NULL is returned. If there is a valid reason not to have this authentication happen, calling `_crypt` avoids authentication.

The `setkey` and `encrypt` entries provide (rather primitive) access to the DES algorithm. The argument of `setkey` is a character array of length 64 containing only the characters with numerical value 0 and 1. If this string is divided into groups of 8, the low-order bit in each group is ignored; this gives a 56-bit key which is set into the machine. This is the key that will be used with the above mentioned algorithm to encrypt or decrypt the string *block* with the function `encrypt`.

The argument to the `encrypt` entry is a character array of length 64 containing only the characters with numerical value 0 and 1. The argument array is modified in place to a similar array representing the bits of the argument after having been subjected to the DES algorithm using the key set by `setkey`. If *edflag* is zero, the argument is encrypted; if non-zero, it is decrypted.

## SEE ALSO

`login(1)`, `passwd(1)`, `getpass(3V)`, `pwdauth(3)`, `passwd(5)`

## BUGS

The return value points to static data whose content is overwritten by each call.

**NAME**

**ctermid** – generate filename for terminal

**SYNOPSIS**

```
#include <stdio.h>
char *ctermid (s)
char *s;
```

**DESCRIPTION**

**ctermid()** generates the pathname of the controlling terminal for the current process, and stores it in a string.

If *s* is a NULL pointer, the string is stored in an internal static area, the contents of which are overwritten at the next call to **ctermid()**, and the address of which is returned. Otherwise, *s* is assumed to point to a character array of at least **L\_ctermid** elements; the path name is placed in this array and the value of *s* is returned. The constant **L\_ctermid** is defined in **<stdio.h>** header file.

**ctermid()** returns a pointer to a null string if it fails, or if the pathname that would refer to the controlling terminal cannot be determined.

**SEE ALSO**

**ttyname(3V)**

**NOTES**

The difference between **ctermid()** and **ttyname(3V)** is that **ttyname()** must be passed a file descriptor and returns the actual name of the terminal associated with that file descriptor, while **ctermid()** returns a string (**/dev/tty**) that will refer to the terminal if used as a file name. Thus **ttyname()** is useful only if the process already has at least one file open to a terminal. **ctermid()** is useful largely for making code portable to (non-UNIX) systems where the current terminal is referred to by a name other than **/dev/tty**.

## NAME

`ctime`, `asctime`, `dysize`, `gmtime`, `localtime`, `strftime`, `strptime`, `timegm`, `timelocal`, `tzset`, `tzsetwall` – convert date and time

## SYNOPSIS

```
#include <time.h>

char *ctime(clock)
time_t *clock;

char *asctime(tm)
struct tm *tm;

int dysize(y)
int y;

struct tm *gmtime(clock)
time_t *clock;

struct tm *localtime(clock)
time_t *clock;

int strftime(buf, bufsize, fmt, tm)
char *buf;
int bufsize;
char *fmt;
struct tm *tm;

char *strptime(buf, fmt, tm)
char *buf;
char *fmt;
struct tm *tm;

time_t timegm(tm)
struct tm *tm;

time_t timelocal(tm)
struct tm *tm;

void tzset()

void tzsetwall()
```

## SYSTEM V SYNOPSIS

In addition to the routines above, the following variables are available:

```
extern long timezone;
extern int daylight;
extern char *tzname[2];
```

## DESCRIPTION

`ctime()` converts a long integer, pointed to by `clock`, to a 26-character string of the form produced by `asctime()`. It first breaks down `clock` to a `tm` structure by calling `localtime()`, and then calls `asctime()` to convert that `tm` structure to a string.

`asctime()` converts a time value contained in a `tm` structure to a 26-character string of the form:

```
Sun Sep 16 01:03:52 1973\n\0
```

Each field has a constant width. `asctime()` returns a pointer to the string.

`dysize()` returns the number of days in the argument year, either 365 or 366. `localtime()` and `gmtime()` return pointers to structures containing the time, broken down into various components of that time represented in a particular time zone. `localtime()` breaks down a time specified by the value pointed to by

the *clock* argument, correcting for the time zone and any time zone adjustments (such as Daylight Savings Time). Before doing so, `localtime()` calls `tzset()` (if `tzset()` has not been called in the current process). `gmtime()` breaks down a time specified by the value pointed to by the *clock* argument into GMT, which is the time the system uses.

`strftime()` converts a time value contained in the `tm` structure pointed to by *tm* to a character string in a format specified by *fmt*. The character string is placed into the array pointed to by *buf*, which is assumed to contain room for at least *buflen* characters. If the result contains no more than *buflen* characters, `strftime()` returns the number of characters produced (not including the terminating null character). Otherwise, it returns zero and the contents of the array are indeterminate. *fmt* is a character string that consists of field descriptors and text characters, reminiscent of `printf(3V)`. Each field descriptor consists of a `%` character followed by another character that specifies the replacement for the field descriptor. All other characters are copied from *fmt* into the result. The following field descriptors are supported:

<code>%%</code>	same as <code>%</code>
<code>%a</code>	day of week, using locale's abbreviated weekday names
<code>%A</code>	day of week, using locale's full weekday names
<code>%b</code>	
<code>%h</code>	month, using locale's abbreviated month names
<code>%B</code>	month, using locale's full month names
<code>%c</code>	date and time as <code>%x %X</code>
<code>%C</code>	date and time, in locale's long-format date and time representation
<code>%d</code>	day of month (01-31)
<code>%D</code>	date as <code>%m/%d/%y</code>
<code>%e</code>	day of month (1-31; single digits are preceded by a blank)
<code>%H</code>	hour (00-23)
<code>%I</code>	hour (00-12)
<code>%j</code>	day number of year (001-366)
<code>%k</code>	hour (0-23; single digits are preceded by a blank)
<code>%l</code>	hour (1-12; single digits are preceded by a blank)
<code>%m</code>	month number (01-12)
<code>%M</code>	minute (00-59)
<code>%n</code>	same as <code>\n</code>
<code>%p</code>	locale's equivalent of AM or PM, whichever is appropriate
<code>%r</code>	time as <code>%I:%M:%S %p</code>
<code>%R</code>	time as <code>%H:%M</code>
<code>%S</code>	seconds (00-59)
<code>%t</code>	same as <code>\t</code>
<code>%T</code>	time as <code>%H:%M:%S</code>
<code>%U</code>	week number of year (01-52), Sunday is the first day of the week
<code>%w</code>	day of week; Sunday is day 0
<code>%W</code>	week number of year (01-52), Monday is the first day of the week
<code>%x</code>	date, using locale's date format
<code>%X</code>	time, using locale's time format

**%y** year within century (00-99)  
**%Y** year, including century (fore example, 1988)  
**%Z** time zone abbreviation

The difference between **%U** and **%W** lies in which day is counted as the first day of the week. Week number 01 is the first week with four or more January days in it.

**strptime()** converts the character string pointed to by *buf* to a time value, which is stored in the **tm** structure pointed to by *tm*, using the format specified by *fmt*. A pointer to the character following the last character in the string pointed to by *buf* is returned. *fmt* is a character string that consists of field descriptors and text characters, reminiscent of **scanf(3v)**. Each field descriptor consists of a **%** character followed by another character that specifies the replacement for the field descriptor. All other characters are copied from *fmt* into the result. The following field descriptors are supported:

**%%** same as **%**  
**%a**  
**%A** day of week, using locale's weekday names; either the abbreviated or full name may be specified  
**%b**  
**%B**  
**%h** month, using locale's month names; either the abbreviated or full name may be specified  
**%c** date and time as **%x %X**  
**%C** date and time, in locale's long-format date and time representation  
**%d**  
**%e** day of month (1-31; leading zeroes are permitted but not required)  
**%D** date as **%m/%d/%y**  
**%H**  
**%k** hour (0-23; leading zeroes are permitted but not required)  
**%I**  
**%l** hour (0-12; leading zeroes are permitted but not required)  
**%j** day number of year (001-366)  
**%m** month number (1-12; leading zeroes are permitted but not required)  
**%M** minute (0-59; leading zeroes are permitted but not required)  
**%p** locale's equivalent of AM or PM  
**%r** time as **%I:%M:%S %p**  
**%R** time as **%H:%M**  
**%S** seconds (0-59; leading zeroes are permitted but not required)  
**%T** time as **%H:%M:%S**  
**%x** date, using locale's date format  
**%X** time, using locale's time format  
**%y** year within century (0-99; leading zeroes are permitted but not required)  
**%Y** year, including century (for example, 1988)

Case is ignored when matching items such as month or weekday names. The **%M**, **%S**, **%y**, and **%Y** fields are optional; if they would be matched by white space, the match is suppressed and the appropriate field of the **tm** structure pointed to by *tm* is left unchanged. If any of the format items **%d**, **%e**, **%H**, **%k**, **%I**, **%l**, **%m**, **%M**, **%S**, **%y**, or **%Y** are matched, but the string that matches them is followed by white

space, all subsequent items in the format string are skipped up to white space or the end of the format. The net result is that, for example, the format `%m/%d/%y` can be matched by the string `12/31`; the `tm_mon` and `tm_mday` fields of the `tm` structure pointed to by `tm` will be set to 11 and 31, respectively, while the `tm_year` field will be unchanged.

`timelocal()` and `timegm()` convert the time specified by the value pointed to by the `tm` argument to a time value that represents that time expressed as the number of seconds since Jan. 1, 1970, 00:00, Greenwich Mean Time. `timelocal()` converts a `tm` structure that represents local time, correcting for the time zone and any time zone adjustments (such as Daylight Savings Time). Before doing so, `timelocal()` calls `tzset()` (if `tzset()` has not been called in the current process). `timegm()` converts a `tm` structure that represents GMT.

`tzset()` uses the value of the environment variable `TZ` to set time conversion information used by `localtime()`. If `TZ` is absent from the environment, the an available approximation to local wall clock time is used by `localtime()`. If `TZ` appears in the environment but its value is a null string, Greenwich Mean Time is used; if `TZ` appears and begins with a slash, it is used as the absolute pathname of the `tzfile-format` (see `tzfile(5)`) file from which to read the time conversion information; if `TZ` appears and begins with a character other than a slash, it is used as a pathname relative to a system time conversion information directory.

`tzsetwall()` sets things up so that `localtime()` returns the best available approximation of local wall clock time.

Declarations of all the functions and externals, and the `tm` structure, are in the `<time.h>` header file. The structure (of type) `tm` structure includes the following fields:

```
int tm_sec;      /* seconds (0 - 59) */
int tm_min;      /* minutes (0 - 59) */
int tm_hour;     /* hours (0 - 23) */
int tm_mday;     /* day of month (1 - 31) */
int tm_mon;      /* month of year (0 - 11) */
int tm_year;     /* year - 1900 */
int tm_wday;     /* day of week (Sunday = 0) */
int tm_yday;     /* day of year (0 - 365) */
int tm_isdst;    /* 1 if DST in effect */
char *tm_zone;   /* abbreviation of timezone name */
long tm_gmtoff;  /* offset from GMT in seconds */
```

`tm_isdst` is non-zero if Daylight Savings Time is in effect. `tm_zone` points to a string that is the name used for the local time zone at the time being converted. `tm_gmtoff` is the offset (in seconds) of the time represented from GMT, with positive values indicating East of Greenwich.

#### SYSTEM V DESCRIPTION

The external `long` variable `timezone` contains the difference, in seconds, between GMT and local standard time (in PST, `timezone` is `8*60*60`). If this difference is not a constant, `timezone` will contain the value of the offset on January 1, 1970 at 00:00 GMT. Since this is not necessarily the same as the value at some particular time, the time in question should be converted to a `tm` structure using `localtime()` and the `tm_gmtoff` field of that structure should be used. The external variable `daylight` is non-zero if and only if Daylight Savings Time would be in effect within the current time zone at some time; it does not indicate whether Daylight Savings Time is currently in effect.

The external variable **tzname** is an array of two **char \*** pointers. The first pointer points to a character string that is the name of the current time zone when Daylight Savings Time is not in effect; the second one, if Daylight Savings Time conversion should be applied, points to a character string that is the name of the current time zone when Daylight Savings Time is in effect. These strings are updated by **localtime()** whenever a time is converted. If Daylight Savings Time is in effect at the time being converted, the second pointer is set to point to the name of the current time zone at that time, otherwise the first pointer is so set.

**timezone**, **daylight**, and **tzname** are retained for compatibility with existing programs.

**FILES**

**/usr/share/lib/zoneinfo**                standard time conversion information directory  
**/usr/share/lib/zoneinfo/localtime**    local time zone file

**SEE ALSO**

**gettimeofday(2)**, **getenv(3V)**, **time(3V)**, **environ(5V)**, **tzfile(5)**

**BUGS**

The return values point to static data, whose contents are overwritten by each call. The **tm\_zone** field of a returned **tm** structure points to a static array of characters, which will also be overwritten at the next call (and by calls to **tzset()** or **tzsetwall()**).

**NAME**

`cctype`, `conv`, `isalpha`, `isupper`, `islower`, `isdigit`, `isxdigit`, `isalnum`, `isspace`, `ispunct`, `isprint`, `isctrl`, `isascii`, `isgraph`, `toupper`, `tolower`, `toascii` – character classification and conversion macros and functions

**SYNOPSIS**

```
#include <cctype.h>
```

```
isalpha(c)
```

```
...
```

**DESCRIPTION****Character Classification Macros**

These macros classify character-coded integer values according to the rules of the coded character set defined by the character type information in the program's locale (category `LC_CTYPE`). On program startup the `LC_CTYPE` category of locale is equivalent to the "C" locale.

In the "C" locale, or in a locale where the character type information is not defined, characters are classified according to the rules of the US-ASCII 7-bit coded character set. The control characters are those below 040 (and the single byte 0177) (DEL). See `ascii(7)`.

In all cases that argument is an `int`, the value of which must be representable as an **unsigned char** or must equal the value of the macro `EOF`. If the argument has any other value, the behavior is undefined.

Each is a predicate returning nonzero for true, zero for false. `isascii()` is defined on all integer values.

<code>isalpha(c)</code>	<code>c</code> is a letter.
<code>isupper(c)</code>	<code>c</code> is an upper case letter.
<code>islower(c)</code>	<code>c</code> is a lower case letter.
<code>isdigit(c)</code>	<code>c</code> is a digit [0-9].
<code>isxdigit(c)</code>	<code>c</code> is a hexadecimal digit [0-9], [A-F], or [a-f].
<code>isalnum(c)</code>	<code>c</code> is an alphanumeric character, that is, <code>c</code> is a letter or a digit.
<code>isspace(c)</code>	<code>c</code> is a SPACE, TAB, RETURN, NEWLINE, FORMFEED, or vertical tab character.
<code>ispunct(c)</code>	<code>c</code> is a punctuation character (neither control nor alphanumeric).
<code>isprint(c)</code>	<code>c</code> is a printing character.
<code>isctrl(c)</code>	<code>c</code> is a delete character or ordinary control character.
<code>isascii(c)</code>	<code>c</code> is an ASCII character, code less than 0200.
<code>isgraph(c)</code>	<code>c</code> is a visible graphic character.

**Character Conversion Macros**

```
toascii(c)
```

Masks `c` with the correct value so that `c` is guaranteed to be an ASCII character in the range 0 through 0x7f. Will not perform mapping from a non-ASCII coded character set into ASCII.

**Character Conversion Functions**

These functions perform simple conversions on single characters. They replace the previous macro definitions which did not extend to support variant settings of the `LC_CTYPE` locale category.

`toupper(c)` Converts `c` to its upper-case equivalent. This function works correctly for all coded character sets and all characters within such sets selected by a valid setting of the `LC_CTYPE` locale category.



**tolower(*c*)** Converts *c* to its lower-case equivalent. This function works correctly for all coded character sets and all characters within such sets selected by a valid setting of the LC\_CTYPE locale category.

If the argument to any of these macros is not in the domain of the function, the result is undefined.

#### SYSTEM V DESCRIPTION

##### Character Conversion Macros

The macros **\_toupper()** and **\_tolower()** are faster than the equivalent functions (**toupper()** and **tolower()**) but only work properly on a restricted range of characters, and will not work on a LC\_CTYPE category other than the default "C" (ASCII).

These macros perform simple conversions on single characters.

**\_toupper(*c*)** converts *c* to its upper-case equivalent. Note: This *only* works where *c* is known to be a lower-case character to start with (presumably checked using **islower()**).

**\_tolower(*c*)** converts *c* to its lower-case equivalent. Note: This *only* works where *c* is known to be an upper-case character to start with (presumably checked using **isupper()**).

#### SEE ALSO

**setlocale(3V)**, **ascii(7)**, **iso\_8859\_1(7)**

**NAME**

curSES – System V terminal screen handling and optimization package

**SYNOPSIS**

The curSES manual page is organized as follows:

In SYNOPSIS

- compiling information
- summary of parameters used by curSES routines

In SYSTEM V SYNOPSIS:

- compiling information

In DESCRIPTION and SYSTEM V DESCRIPTION:

- An overview of how curSES routines should be used

In ROUTINES, descriptions of curSES routines are grouped under the appropriate topics:

- Overall Screen Manipulation
- Window and Pad Manipulation
- Output
- Input
- Output Options Setting
- Input Options Setting
- Environment Queries
- Low-level Curses Access
- Miscellaneous
- Use of `curscr`

In SYSTEM V ROUTINES, descriptions of curSES routines are grouped under the appropriate topics:

- Overall Screen Manipulation
- Window and Pad Manipulation
- Output
- Input
- Output Options Setting
- Input Options Setting
- Environment Queries
- Soft Labels
- Low-level Curses Access
- TermInfo-Level Manipulations
- Termcap Emulation
- Miscellaneous
- Use of `curscr`

Then come sections on:

- SYSTEM V ATTRIBUTES
- SYSTEM V FUNCTION KEYS

- LINE GRAPHICS

`cc [ flags ] files -lcurses -ltermcap [ libraries ]`

`#include <curses.h>` (automatically includes `<stdio.h>` and `<unctl.h>`.)

The parameters in the following list are not global variables. This is a summary of the parameters used by the curses library routines. All routines return the `int` values `ERR` or `OK` unless otherwise noted. Routines that return pointers always return `NULL` on error. `ERR`, `OK`, and `NULL` are all defined in `<curses.h>`.) Routines that return integers are not listed in the parameter list below.

**bool bf**

**char \*\*area, \*boolnames[ ], \*boolcodes[ ], \*boolfnames[ ], \*bp  
char \*cap, \*capname, codename[2], erasechar, \*filename, \*fmt**

**char \*keyname, killchar, \*label, \*longname**

**char \*name, \*numnames[ ], \*numcodes[ ], \*numfnames[ ]**

**char \*slk\_label, \*str, \*strnames[ ], \*strcodes[ ], \*strfnames[ ]**

**char \*term, \*tgetstr, \*tigetstr, \*tgoto, \*tparm, \*type**

**chtype attrs, ch, horch, vertch**

**FILE \*infd, \*outfd**

**int begin\_x, begin\_y, begline, bot, c, col, count**

**int dmaxcol, dmaxrow, dmincol, dminrow, \*errret, fildes**

**int (\*init()), labfmt, labnum, line**

**int ms, ncols, new, newcol, newrow, nlines, numlines**

**int oldcol, oldrow, overlay**

**int p1, p2, p9, pmincol, pminrow, (\*putc()), row**

**int smaxcol, smaxrow, smincol, sminrow, start**

**int tenths, top, visibility, x, y**

**SCREEN \*new, \*newterm, \*set\_term**

**TERMINAL \*cur\_term, \*nterm, \*oterm**

**va\_list varglist**

**WINDOW \*curscr, \*dstwin, \*initscr, \*newpad, \*newwin, \*orig**

**WINDOW \*pad, \*srcwin, \*stdscr, \*subpad, \*subwin, \*win**

#### SYSTEM V SYNOPSIS

`/usr/5bin/cc [ flag ... ] file ... -lcurses [ library ... ]`

`#include <curses.h>` (automatically includes `<stdio.h>`, `<termio.h>`, and `<unctrl.h>`.)

#### DESCRIPTION

These routines give the user a method of updating screens with reasonable optimization. They keep an image of the current screen, and the user sets up an image of a new one. Then the `refresh()` tells the routines to make the current screen look like the new one. In order to initialize the routines, the routine `initscr()` must be called before any of the other routines that deal with windows and screens are used. The routine `endwin()` should be called before exiting.

#### SYSTEM V DESCRIPTION

The `curses` routines give the user a terminal-independent method of updating screens with reasonable optimization.

In order to initialize the routines, the routine `initscr()` or `newterm()` must be called before any of the other routines that deal with windows and screens are used. Three exceptions are noted where they apply. The routine `endwin()` must be called before exiting. To get character-at-a-time input without echoing, (most interactive, screen oriented programs want this) after calling `initscr()` you should call `'cbreak (); noecho ();'` Most programs would additionally call `'nonl (); intrflush(stdscr, FALSE); keypad(stdscr, TRUE);'`.

Before a **curses** program is run, a terminal's TAB stops should be set and its initialization strings, if defined, must be output. This can be done by executing the **tset** command in your **.profile** or **.login** file. For further details, see **tset(1)** and the **Tabs and Initialization** subsection of **terminfo(5V)**.

The **curses** library contains routines that manipulate data structures called *windows* that can be thought of as two-dimensional arrays of characters representing all or part of a terminal screen. A default window called **stdscr** is supplied, which is the size of the terminal screen. Others may be created with **newwin()**. Windows are referred to by variables declared as **WINDOW \***; the type **WINDOW** is defined in **<curses.h>** to be a C structure. These data structures are manipulated with routines described below, among which the most basic are **move()** and **addch()**. More general versions of these routines are included with names beginning with **w**, allowing you to specify a window. The routines not beginning with **w** usually affect **stdscr**. Then **refresh()** is called, telling the routines to make the user's terminal screen look like **stdscr**. The characters in a window are actually of type **chtype**, so that other information about the character may also be stored with each character.

Special windows called *pads* may also be manipulated. These are windows that are not constrained to the size of the screen and whose contents need not be displayed completely. See the description of **newpad()** under **Window and Pad Manipulation** for more information.

In addition to drawing characters on the screen, video attributes may be included that cause the characters to show up in modes such as underlined or in reverse video on terminals that support such display enhancements. Line drawing characters may be specified to be output. On input, **curses** is also able to translate arrow and function keys that transmit escape sequences into single values. The video attributes, line drawing characters, and input values use names, defined in **<curses.h>**, such as **A\_REVERSE**, **ACS\_HLINE**, and **KEY\_LEFT**.

**curses** also defines the **WINDOW \*** variable, **curscr**, which is used only for certain low-level operations like clearing and redrawing a garbaged screen. **curscr** can be used in only a few routines. If the window argument to **clearok()** is **curscr**, the next call to **wrefresh()** with any window will clear and repaint the screen from scratch. If the window argument to **wrefresh()** is **curscr**, the screen is immediately cleared and repainted from scratch. This is how most programs would implement a "repaint-screen" function. More information on using **curscr** is provided where its use is appropriate.

The environment variables **LINES** and **COLUMNS** may be set to override **curses**'s idea of how large a screen is.

If the environment variable **TERMINFO** is defined, any program using **curses** will check for a local terminal definition before checking in the standard place. For example, if the environment variable **TERM** is set to **sun**, then the compiled terminal definition is found in **/usr/share/lib/terminfo/s/sun**. The **s** is copied from the first letter of **sun** to avoid creation of huge directories.) However, if **TERMINFO** is set to **\$HOME/myterms**, **curses** will first check **\$HOME/myterms/s/sun**, and, if that fails, will then check **/usr/share/lib/terminfo/s/sun**. This is useful for developing experimental definitions or when write permission on **/usr/share/lib/terminfo** is not available.

The integer variables **LINES** and **COLS** are defined in **<curses.h>**, and will be filled in by **initscr()** with the size of the screen. For more information, see the subsection **Terminfo-Level Manipulations**. The constants **TRUE** and **FALSE** have the values **1** and **0**, respectively. The constants **ERR** and **OK** are returned by routines to indicate whether the routine successfully completed. These constants are also defined in **<curses.h>**.

## ROUTINES

Many of the following routines have two or more versions. The routines prefixed with **w** require a *window* argument. The routines prefixed with **p** require a *pad* argument. Those without a prefix generally use **stdscr**.

The routines prefixed with **mv** require *y* and *x* coordinates to move to before performing the appropriate action. The **mv** routines imply a call to **move()** before the call to the other routine. The window argument is always specified before the coordinates. *y* always refers to the row (of the window), and *x* always refers to the column. The upper left corner is always (0,0), not (1,1). The routines prefixed with **mvw** take both a *window* argument and *y* and *x* coordinates.

In each case, *win* is the window affected and *pad* is the pad affected. (*win* and *pad* are always of type WINDOW \*.) Option-setting routines require a boolean flag *bf* with the value TRUE or FALSE. (*bf* is always of type bool.) The types WINDOW, bool, and chtype are defined in < curses.h > (see SYNOPSIS for a summary of what types all variables are).

All routines return either the integer ERR or the integer OK, unless otherwise noted. Routines that return pointers always return NULL on error.

#### Overall Screen Manipulation

**WINDOW \*initscr()** The first routine called should almost always be **initscr()**. The exceptions are **slk\_init()**, **filter()**, and **ripline()**. This will determine the terminal type and initialize all curses data structures. **initscr()** also arranges that the first call to **refresh()** will clear the screen. If errors occur, **initscr()** will write an appropriate error message to standard error and exit; otherwise, a pointer to **stdscr** is returned. If the program wants an indication of error conditions, **newterm()** should be used instead of **initscr()**. **initscr()** should only be called once per application.

**endwin()** A program should always call **endwin()** before exiting or escaping from curses mode temporarily, to do a shell escape or **system(3)** call, for example. This routine will restore **termio(4)** modes, move the cursor to the lower left corner of the screen and reset the terminal into the proper non-visual mode. To resume after a temporary escape, call **wrefresh()** or **doupdate()**.

#### Window and Pad Manipulation

**refresh()**

**wrefresh(win)** These routines (or **prefresh()**, **pnoutrefresh()**, **wnoutrefresh()**, or **doupdate()**) must be called to write output to the terminal, as most other routines merely manipulate data structures. **wrefresh()** copies the named window to the physical terminal screen, taking into account what is already there in order to minimize the amount of information that's sent to the terminal (called optimization). **refresh()** does the same thing, except it uses **stdscr** as a default window. Unless **leaveok()** has been enabled, the physical cursor of the terminal is left at the location of the window's cursor. The number of characters output to the terminal is returned.

Note: **refresh()** is a macro.

**WINDOW \*newwin(nlines, ncols, begin\_y, begin\_x)**

Create and return a pointer to a new window with the given number of lines (or rows), *nlines*, and columns, *ncols*. The upper left corner of the window is at line *begin\_y*, column *begin\_x*. If either *nlines* or *ncols* is 0, they will be set to the value of *lines-begin\_y* and *cols-begin\_x*. A new full-screen window is created by calling **newwin(0,0,0,0)**.

**mvwin(win, y, x)**

Move the window so that the upper left corner will be at position (*y*, *x*). If the move would cause the window to be off the screen, it is an error and the window is not moved.

**WINDOW \*subwin(orig, nlines, ncols, begin\_y, begin\_x)**

Create and return a pointer to a new window with the given number of lines (or rows), *nlines*, and columns, *ncols*. The window is at position (*begin\_y*, *begin\_x*) on the screen. This position is relative to the screen, and not to the window *orig*. The window is made in the middle of the window *orig*, so that changes made to

one window will affect both windows. When using this routine, often it will be necessary to call `touchwin()` or `touchline()` on *orig* before calling `wrefresh`.

**delwin** (*win*)

Delete the named window, freeing up all memory associated with it. In the case of overlapping windows, subwindows should be deleted before the main window.

#### Output

These routines are used to “draw” text on windows.

**addch** (*ch*)

**waddch** (*win, ch*)

**mvaddch** (*y, x, ch*)

**mvwaddch** (*win, y, x, ch*)

The character *ch* is put into the window at the current cursor position of the window and the position of the window cursor is advanced. Its function is similar to that of `putchar()` (see `putc(3s)`). At the right margin, an automatic newline is performed. At the bottom of the scrolling region, if `scrollok()` is enabled, the scrolling region will be scrolled up one line.

If *ch* is a TAB, NEWLINE, or backspace, the cursor will be moved appropriately within the window. A NEWLINE also does a `clrtoeol()` before moving. TAB characters are considered to be at every eighth column. If *ch* is another control character, it will be drawn in the CTRL-X notation. (Calling `winch()` after adding a control character will not return the control character, but instead will return the representation of the control character.)

Video attributes can be combined with a character by or-ing them into the parameter. This will result in these attributes also being set. The intent here is that text, including attributes, can be copied from one place to another using `inch()` and `addch()`. See `standout()`, below.

Note: *ch* is actually of type `chtype`, not a character.

Note: `addch()`, `mvaddch()`, and `mvwaddch()` are macros.

**addstr** (*str*)

**waddstr** (*win, str*)

**mvwaddstr** (*win, y, x, str*)

**mvaddstr** (*y, x, str*)

These routines write all the characters of the null-terminated character string *str* on the given window. This is equivalent to calling `waddch()` once for each character in the string.

Note: `addstr()`, `mvaddstr()`, and `mvwaddstr()` are macros.

**box** (*win, vertch, horch*)

A box is drawn around the edge of the window, *win*. *vertch* and *horch* are the characters the box is to be drawn with. If *vertch* and *horch* are 0, then appropriate default characters, `ACS_VLINE` and `ACS_HLINE`, will be used.

Note: *vertch* and *horch* are actually of type `chtype`, not characters.

**erase()**

**werase** (*win*)

These routines copy blanks to every position in the window.

Note: `erase()` is a macro.

**clear()****wclear** (*win*)

These routines are like **erase()** and **werase()**, but they also call **clearok()**, arranging that the screen will be cleared completely on the next call to **wrefresh()** for that window, and repainted from scratch.

Note: **clear()** is a macro.

**clrtobot()****wclrtobot** (*win*)

All lines below the cursor in this window are erased. Also, the current line to the right of the cursor, inclusive, is erased.

Note: **clrtobot()** is a macro.

**clrtoeol()****wclrtoeol** (*win*)

The current line to the right of the cursor, inclusive, is erased.

Note: **clrtoeol()** is a macro.

**delch()****wdelch** (*win*)**mvdelch** (*y, x*)**mvwdelch** (*win, y, x*)

The character under the cursor in the window is deleted. All characters to the right on the same line are moved to the left one position and the last character on the line is filled with a blank. The cursor position does not change (after moving to (*y, x*), if specified). This does not imply use of the hardware "delete-character" feature.

Note: **delch()**, **mvdelch()**, and **mvwdelch()** are macros.

**deleteln()****wdeleteln** (*win*)

The line under the cursor in the window is deleted. All lines below the current line are moved up one line. The bottom line of the window is cleared. The cursor position does not change. This does not imply use of the hardware "delete-line" feature.

Note: **deleteln()** is a macro.

**getyx** (*win, y, x*)

The cursor position of the window is placed in the two integer variables *y* and *x*. This is implemented as a macro, so no '&' is necessary before the variables.

**insch** (*ch*)**winsch** (*win, ch*)**mvwinsch** (*win, y, x, ch*)**mvinsch** (*y, x, ch*)

The character *ch* is inserted before the character under the cursor. All characters to the right are moved one SPACE to the right, possibly losing the rightmost character of the line. The cursor position does not change (after moving to (*y, x*), if specified). This does not imply use of the hardware "insert-character" feature.

Note: *ch* is actually of type **chtype**, not a character.

Note: **insch()**, **mvinsch()**, and **mvwinsch()** are macros.

**insertln()****winsertln** (*win*)

A blank line is inserted above the current line and the bottom line is lost. This does not imply use of the hardware "insert-line" feature.

Note: **insertln()** is a macro.

**move** (*y, x*)

**wmove** (*win, y, x*) The cursor associated with the window is moved to line (row) *y*, column *x*. This does not move the physical cursor of the terminal until **refresh()** is called. The position specified is relative to the upper left corner of the window, which is (0, 0).

Note: **move()** is a macro.

**overlay** (*srcwin, dstwin*)

**overwrite** (*srcwin, dstwin*)

These routines overlay *srcwin* on top of *dstwin*; that is, all text in *srcwin* is copied into *dstwin*. *srcwin* and *dstwin* need not be the same size; only text where the two windows overlap is copied. The difference is that **overlay()** is non-destructive (blanks are not copied), while **overwrite()** is destructive.

**printw** (*fmt* [, *arg...*])

**wprintw** (*win, fmt* [, *arg...*])

**mvprintw** (*y, x, fmt* [, *arg...*])

**mvwprintw** (*win, y, x, fmt* [, *arg...*])

These routines are analogous to **printf(3V)**. The string that would be output by **printf(3V)** is instead output using **waddstr()** on the given window.

**scroll** (*win*)

The window is scrolled up one line. This involves moving the lines in the window data structure. As an optimization, if the window is **stdscr** and the scrolling region is the entire window, the physical screen will be scrolled at the same time.

**touchwin** (*win*)

**touchline** (*win, start, count*)

Throw away all optimization information about which parts of the window have been touched, by pretending that the entire window has been drawn on. This is sometimes necessary when using overlapping windows, since a change to one window will affect the other window, but the records of which lines have been changed in the other window will not reflect the change. **touchline()** only pretends that *count* lines have been changed, beginning with line *start*.

#### Input

**getch()**

**wgetch** (*win*)

**mvgetch** (*y, x*)

**mvwgetch** (*win, y, x*)

A character is read from the terminal associated with the window. In **NODELAY** mode, if there is no input waiting, the value **ERR** is returned. In **DELAY** mode, the program will hang until the system passes text through to the program. Depending on the setting of **cbreak()**, this will be after one character (**CBREAK** mode), or after the first newline (**NOCBREAK** mode). In **HALF-DELAY** mode, the program will hang until a character is typed or the specified timeout has been reached. Unless **noecho()** has been set, the character will also be echoed into the designated window. No **refresh()** will occur between the **move()** and the **getch()** done within the routines **mvgetch()** and **mvwgetch()**.

When using **getch()**, **wgetch()**, **mvgetch()**, or **mvwgetch()**, do not set both **NOCBREAK** mode (**nocbreak()**) and **ECHO** mode (**echo()**) at the same time. Depending on the state of the terminal driver when each character is typed, the program may produce undesirable results.



If `keypad` (*win*, TRUE) has been called, and a function key is pressed, the token for that function key will be returned instead of the raw characters. See `keypad()` under **Input Options Setting**. Possible function keys are defined in `< curses.h >` with integers beginning with 0401, whose names begin with `KEY_`. If a character is received that could be the beginning of a function key (such as escape), `curses` will set a timer. If the remainder of the sequence is not received within the designated time, the character will be passed through, otherwise the function key value will be returned. For this reason, on many terminals, there will be a delay after a user presses the escape key before the escape is returned to the program. Use by a programmer of the escape key for a single character routine is discouraged. Also see `notimeout()` below.

Note: `getch()`, `mvgetch()`, and `mvwgetch()` are macros.

`getstr` (*str*)

`wgetstr` (*win*, *str*)

`mvgetstr` (*y*, *x*, *str*)

`mvwgetstr` (*win*, *y*, *x*, *str*)

A series of calls to `getch()` is made, until a newline, carriage return, or enter key is received. The resulting value is placed in the area pointed at by the character pointer *str*. The user's erase and kill characters are interpreted. As in `mvgetch()`, no `refresh()` is done between the `move()` and `getstr()` within the routines `mvgetstr()` and `mvwgetstr()`.

Note: `getstr()`, `mvgetstr()`, and `mvwgetstr()` are macros.

`inch()`

`winch` (*win*)

`mvinch` (*y*, *x*)

`mvwinch` (*win*, *y*, *x*)

The character, of type `chtype`, at the current position in the named window is returned. If any attributes are set for that position, their values will be OR'ed into the value returned. The predefined constants `A_CHARTEXT` and `A_ATTRIBUTES`, defined in `< curses.h >`, can be used with the C logical AND (&) operator to extract the character or attributes alone.

Note: `inch()`, `winch()`, `mvinch()`, and `mvwinch()` are macros.

`scanw` (*fmt* [, *arg* ... ] )

`wscanw` (*win*, *fmt* [, *arg* ... ])

`mvscanw` (*y*, *x*, *fmt* [, *arg* ... ])

`mvwscanw` (*win*, *y*, *x*, *fmt* [, *arg* ... ])

These routines correspond to `scanf(3V)`, as do their arguments and return values. `wgetstr()` is called on the window, and the resulting line is used as input for the scan.

#### Output Options Setting

These routines set options within `curses` that deal with output. All options are initially FALSE, unless otherwise stated. It is not necessary to turn these options off before calling `endwin()`.

`clearok` (*win*, *bf*)

If enabled (*bf* is TRUE), the next call to `wrefresh()` with this window will clear the screen completely and redraw the entire screen from scratch. This is useful when the contents of the screen are uncertain, or in some cases for a more pleasing visual effect.

- idlok** (*win, bf*) If enabled (*bf* is **TRUE**), **curses** will consider using the hardware “insert/delete-line” feature of terminals so equipped. If disabled (*bf* is **FALSE**), **curses** will very seldom use this feature. The “insert/delete-character” feature is always considered. This option should be enabled only if your application needs “insert/delete-line”, for example, for a screen editor. It is disabled by default because “insert/delete-line” tends to be visually annoying when used in applications where it is not really needed. If “insert/delete-line” cannot be used, **curses** will redraw the changed portions of all lines.
- leaveok** (*win, bf*) Normally, the hardware cursor is left at the location of the window cursor being refreshed. This option allows the cursor to be left wherever the update happens to leave it. It is useful for applications where the cursor is not used, since it reduces the need for cursor motions. If possible, the cursor is made invisible when this option is enabled.
- scrollok** (*win, bf*) This option controls what happens when the cursor of a window is moved off the edge of the window or scrolling region, either from a newline on the bottom line, or typing the last character of the last line. If disabled (*bf* is **FALSE**), the cursor is left on the bottom line at the location where the offending character was entered. If enabled (*bf* is **TRUE**), **wrefresh()** is called on the window, and then the physical terminal and window are scrolled up one line. Note: in order to get the physical scrolling effect on the terminal, it is also necessary to call **idlok()**.

**nl()**

**nonl()**

These routines control whether **NEWLINE** is translated into **RETURN** and **LINEFEED** on output, and whether **RETURN** is translated into **NEWLINE** on input. Initially, the translations do occur. By disabling these translations using **nonl()**, **curses** is able to make better use of the linefeed capability, resulting in faster cursor motion.

#### Input Options Setting

These routines set options within **curses** that deal with input. The options involve using **ioctl(2)** and therefore interact with **curses** routines. It is not necessary to turn these options off before calling **endwin()**.

For more information on these options, refer to *Programming Utilities and Libraries*.

**cbreak()**

**nocbreak()**

These two routines put the terminal into and out of **CBREAK** mode, respectively. In **CBREAK** mode, characters typed by the user are immediately available to the program and erase/kill character processing is not performed. When in **NOCBREAK** mode, the tty driver will buffer characters typed until a **NEWLINE** or **RETURN** is typed. Interrupt and flow-control characters are unaffected by this mode (see **termio(4)**). Initially the terminal may or may not be in **CBREAK** mode, as it is inherited, therefore, a program should call **cbreak()** or **nocbreak()** explicitly. Most interactive programs using **curses** will set **CBREAK** mode.

Note: **cbreak()** overrides **raw()**. See **getch()** under **Input** for a discussion of how these routines interact with **echo()** and **noecho()**.

**echo()**

**noecho()**

These routines control whether characters typed by the user are echoed by **getch()** as they are typed. Echoing by the tty driver is always disabled, but initially **getch()** is in **ECHO** mode, so characters typed are echoed. Authors of most interactive programs prefer to do their own echoing in a controlled area of the screen, or not to echo at all, so they disable echoing by calling **noecho()**. See **getch()** under **Input** for a discussion of how these routines interact with **cbreak()** and **nocbreak()**.

**raw()**

**noraw()**

The terminal is placed into or out of RAW mode. RAW mode is similar to CBREAK mode, in that characters typed are immediately passed through to the user program. The differences are that in RAW mode, the interrupt, quit, suspend, and flow control characters are passed through uninterpreted, instead of generating a signal. RAW mode also causes 8-bit input and output. The behavior of the BREAK key depends on other bits in the terminal driver that are not set by `curses`.

#### Environment Queries

**baudrate()**

Returns the output speed of the terminal. The number returned is in bits per second, for example, 9600, and is an integer.

**char erasechar()**

The user's current erase character is returned.

**char killchar()**

The user's current line-kill character is returned.

**char \*longname()**

This routine returns a pointer to a static area containing a verbose description of the current terminal. The maximum length of a verbose description is 128 characters. It is defined only after the call to `initscr()` or `newterm()`. The area is overwritten by each call to `newterm()` and is not restored by `set_term()`, so the value should be saved between calls to `newterm()` if `longname()` is going to be used with multiple terminals.

#### Low-Level curses Access

The following routines give low-level access to various `curses` functionality. These routines typically would be used inside of library routines.

**resetty()**

**savetty()**

These routines save and restore the state of the terminal modes. `savetty()` saves the current state of the terminal in a buffer and `resetty()` restores the state to what it was at the last call to `savetty()`.

#### Miscellaneous

**unctrl(*c*)**

This macro expands to a character string which is a printable representation of the character *c*. Control characters are displayed in the ^X notation. Printing characters are displayed as is.

`unctrl()` is a macro, defined in `<unctrl.h>`, which is automatically included by `<curses.h>`.

**flusok(*win,boolf*)**

set flush-on-refresh flag for *win*

**getcap(*name*)**

get terminal capability *name*

**touchoverlap(*win1,win2*)**

mark overlap of *win1* on *win2* as changed

#### Use of curscr

The special window `curscr` can be used in only a few routines. If the window argument to `clearok()` is `curscr`, the next call to `wrefresh()` with any window will cause the screen to be cleared and repainted from scratch. If the window argument to `wrefresh()` is `curscr`, the screen is immediately cleared and repainted from scratch. This is how most programs would implement a "repaint-screen" routine. The source window argument to `overlay()`, `overwrite()`, and `copywin` may be `curscr`, in which case the current contents of the virtual terminal screen will be accessed.

#### Obsolete Calls

Various routines are provided to maintain compatibility in programs written for older versions of the `curses` library. These routines are all emulated as indicated below.

**crmode()** Replaced by **cbreak()**.  
**gettmode()** A no-op.  
**nocrmode()** Replaced by **nocbreak()**.

#### SYSTEM V ROUTINES

The above routines are available as described except for **flusok()**, **getcap()** and **touchoverlap()** which are not available.

In addition, the following routines are available:

#### Overall Screen Manipulation

**isendwin()** Returns **TRUE** if **endwin()** has been called without any subsequent calls to **wrefresh()**.

#### SCREEN \*newterm(*type, outfd, infd*)

A program that outputs to more than one terminal must use **newterm()** for each terminal instead of **initscr()**. A program that wants an indication of error conditions, so that it may continue to run in a line-oriented mode if the terminal cannot support a screen-oriented program, must also use this routine. **newterm()** should be called once for each terminal. It returns a variable of type **SCREEN\*** that should be saved as a reference to that terminal. The arguments are the *type* of the terminal to be used in place of the environment variable **TERM**; *outfd*, a **stdio(3V)** file pointer for output to the terminal; and *infd*, another file pointer for input from the terminal. When it is done running, the program must also call **endwin()** for each terminal being used. If **newterm()** is called more than once for the same terminal, the first terminal referred to must be the last one for which **endwin()** is called.

#### SCREEN \*set\_term(*new*)

This routine is used to switch between different terminals. The screen reference *new* becomes the new current terminal. A pointer to the screen of the previous terminal is returned by the routine. This is the only routine that manipulates **SCREEN** pointers; all other routines affect only the current terminal.

#### Window and Pad Manipulation

##### wnoutrefresh(*win*)

##### doupdate()

These two routines allow multiple updates to the physical terminal screen with more efficiency than **wrefresh()** alone. How this is accomplished is described in the next paragraph.

**curses** keeps two data structures representing the terminal screen: a *physical* terminal screen, describing what is actually on the screen, and a *virtual* terminal screen, describing what the programmer wants to have on the screen. **wrefresh()** works by first calling **wnoutrefresh()**, which copies the named window to the virtual screen, and then by calling **doupdate()**, which compares the virtual screen to the physical screen and does the actual update. If the programmer wishes to output several windows at once, a series of calls to **wrefresh()** will result in alternating calls to **wnoutrefresh()** and **doupdate()**, causing several bursts of output to the screen. By first calling **wnoutrefresh()** for each window, it is then possible to call **doupdate()** once, resulting in only one burst of output, with probably fewer total characters transmitted and certainly less processor time used.

#### WINDOW \*newpad(*nlines, ncols*)

Create and return a pointer to a new pad data structure with the given number of lines (or rows), *nlines*, and columns, *ncols*. A pad is a window that is not restricted by the screen size and is not necessarily associated with a particular part of the screen. Pads can be used when a large window is needed, and only a part of

the window will be on the screen at one time. Automatic refreshes of pads (for example, from scrolling or echoing of input) do not occur. It is not legal to call **wrefresh()** with a pad as an argument; the routines **prefresh()** or **pnoutrefresh()** should be called instead. Note: these routines require additional parameters to specify the part of the pad to be displayed and the location on the screen to be used for display.

**WINDOW \*subpad** (*orig, nlines, ncols, begin\_y, begin\_x*)

Create and return a pointer to a subwindow within a pad with the given number of lines (or rows), *nlines*, and columns, *ncols*. Unlike **subwin()**, which uses screen coordinates, the window is at position (*begin\_y, begin\_x*) on the pad. The window is made in the middle of the window *orig*, so that changes made to one window will affect both windows. When using this routine, often it will be necessary to call **touchwin()** or **touchline()** on *orig* before calling **prefresh()**.

**prefresh** (*pad, pminrow, pmincol, sminrow, smincol, smaxrow, smaxcol*)

**pnoutrefresh** (*pad, pminrow, pmincol, sminrow, smincol, smaxrow, smaxcol*)

These routines are analogous to

**wrefresh()** and **wnoutrefresh()** except that pads, instead of windows, are involved. The additional parameters are needed to indicate what part of the pad and screen are involved. *pminrow* and *pmincol* specify the upper left corner, in the pad, of the rectangle to be displayed. *sminrow, smincol, smaxrow, and smaxcol* specify the edges, on the screen, of the rectangle to be displayed in. The lower right corner in the pad of the rectangle to be displayed is calculated from the screen coordinates, since the rectangles must be the same size. Both rectangles must be entirely contained within their respective structures. Negative values of *pminrow, pmincol, sminrow, or smincol* are treated as if they were zero.

#### Output

These routines are used to “draw” text on windows.

**echochar** (*ch*)

**wechochar** (*win, ch*)

**pechochar** (*pad, ch*)

These routines are functionally equivalent to a call to **addch** (*ch*) followed by a call to **refresh()**, a call to **waddch** (*win, ch*) followed by a call to **wrefresh** (*win*), or a call to **waddch** (*pad, ch*) followed by a call to **prefresh** (*pad*). The knowledge that only a single character is being output is taken into consideration and, for non-control characters, a considerable performance gain can be seen by using these routines instead of their equivalents. In the case of **pechochar()**, the last location of the pad on the screen is reused for the arguments to **prefresh()**.

Note: *ch* is actually of type **chtype**, not a character.

Note: **echochar()** is a macro.

**attroff** (*attrs*)  
**wattroff** (*win, attrs*)  
**attron** (*attrs*)  
**wattron** (*win, attrs*)  
**attrset** (*attrs*)  
**wattrset** (*win, attrs*)

**beep**()

**flash**()

These routines are used to signal the terminal user. **beep**() will sound the audible alarm on the terminal, if possible, and if not, will flash the screen (visible bell), if that is possible. **flash**() will flash the screen, and if that is not possible, will sound the audible signal. If neither signal is possible, nothing will happen. Nearly all terminals have an audible signal (bell or beep) but only some can flash the screen.

**delay\_output** (*ms*)

Insert a *ms* millisecond pause in the output. It is not recommended that this routine be used extensively, because padding characters are used rather than a processor pause.

**getbegyx** (*win, y, x*)

**getmaxyx** (*win, y, x*)

Like **getyx**(), these routines store the current beginning coordinates and size of the specified window.

Note: **getbegyx**() and **getmaxyx**() are macros.

**copywin** (*srcwin, dstwin, sminrow, smincol, dminrow, dmincol, dmaxrow, dmaxcol, overlay*)

This routine provides a finer grain of control over the **overlay**() and **overwrite**() routines. Like in the **prefresh**() routine, a rectangle is specified in the destination window, (*dminrow, dmincol*) and (*dmaxrow, dmaxcol*), and the upper-left-corner coordinates of the source window, (*sminrow, smincol*). If the argument *overlay* is true, then copying is non-destructive, as in **overlay**() .

**vwprintw** (*win, fmt, varglist*)

This routine corresponds to **vprintf**(3V). It performs a **printw**() using a variable argument list. The third argument is a *va\_list*, a pointer to a list of arguments, as defined in *<varargs.h>*. See the **vprintf**(3V) and **varargs**(3) manual pages for a detailed description on how to use variable argument lists.

#### Input

**flushinp**()

Throws away any typeahead that has been typed by the user and has not yet been read by the program.

**ungetch** (*c*)

Place *c* back onto the input queue to be returned by the next call to **wgetch**() .

**vwscanw** (*win, fmt, ap*)

This routine is similar to **vwprintw**() above in that performs a **wscanw**() using a variable argument list. The third argument is a *va\_list*, a pointer to a list of arguments, as defined in *<varargs.h>*. See the **vprintf**(3V) and **varargs**(3) manual pages for a detailed description on how to use variable argument lists.

#### Output Options Setting

These routines set options within **curses** that deal with output. All options are initially FALSE, unless otherwise stated. It is not necessary to turn these options off before calling **endwin**() .

**setscreg** (*top, bot*)

**wsetscreg** (*win, top, bot*)

These routines allow the user to set a software scrolling region in a window. *top* and *bot* are the line numbers of the top and bottom margin of the scrolling region. Line 0 is the top line of the window. If this option and **scrollok**() are enabled, an

attempt to move off the bottom margin line will cause all lines in the scrolling region to scroll up one line. Note: this has nothing to do with use of a physical scrolling region capability in the terminal, like that in the DEC VT100. Only the text of the window is scrolled; if `idlok()` is enabled and the terminal has either a scrolling region or “insert/delete-line” capability, they will probably be used by the output routines.

Note: `setscrreg()` and `wsetscrreg()` are macros.

### Input Options Setting

These routines set options within `curses` that deal with input. The options involve using `ioctl(2)` and therefore interact with `curses` routines. It is not necessary to turn these options off before calling `endwin()`.

For more information on these options, refer to *Programming Utilities and Libraries*.

**halfdelay** (*tenths*) Half-delay mode is similar to CBREAK mode in that characters typed by the user are immediately available to the program. However, after blocking for *tenths* tenths of seconds, ERR will be returned if nothing has been typed. *tenths* must be a number between 1 and 255. Use `nocbreak()` to leave half-delay mode.

**intrflush** (*win, bf*) If this option is enabled, when an interrupt key is pressed on the keyboard (interrupt, break, quit) all output in the tty driver queue will be flushed, giving the effect of faster response to the interrupt, but causing `curses` to have the wrong idea of what is on the screen. Disabling the option prevents the flush. The default for the option is inherited from the tty driver settings. The window argument is ignored.

**keypad** (*win, bf*) This option enables the keypad of the user’s terminal. If enabled, the user can press a function key (such as an arrow key) and `wgetch()` will return a single value representing the function key, as in `KEY_LEFT`. If disabled, `curses` will not treat function keys specially and the program would have to interpret the escape sequences itself. If the keypad in the terminal can be turned on (made to transmit) and off (made to work locally), turning on this option will cause the terminal keypad to be turned on when `wgetch()` is called.

**meta** (*win, bf*) If enabled, characters returned by `wgetch()` are transmitted with all 8 bits, instead of with the highest bit stripped. In order for `meta()` to work correctly, the `km` (`has_meta_key`) capability has to be specified in the terminal’s `terminfo(5V)` entry.

**nodelay** (*win, bf*) This option causes `wgetch()` to be a non-blocking call. If no input is ready, `wgetch()` will return ERR. If disabled, `wgetch()` will hang until a key is pressed.

**notimeout** (*win, bf*) While interpreting an input escape sequence, `wgetch()` will set a timer while waiting for the next character. If `notimeout(win, TRUE)` is called, then `wgetch()` will not set a timer. The purpose of the timeout is to differentiate between sequences received from a function key and those typed by a user.

**typeahead** (*fildes*) `curses` does “line-breakout optimization” by looking for typeahead periodically while updating the screen. If input is found, and it is coming from a tty, the current update will be postponed until `refresh()` or `doupdate()` is called again. This allows faster response to commands typed in advance. Normally, the file descriptor for the input FILE pointer passed to `newterm()`, or `stdin` in the case that `initscr()` was used, will be used to do this typeahead checking. The `typeahead()` routine specifies that the file descriptor *fildes* is to be used to check for typeahead instead. If *fildes* is `-1`, then no typeahead checking will be done.

Note: *fildes* is a file descriptor, not a `<stdio.h>` FILE pointer.

**Environment Queries**

- has\_ic()** True if the terminal has insert- and delete-character capabilities.
- has\_il()** True if the terminal has insert- and delete-line capabilities, or can simulate them using scrolling regions. This might be used to check to see if it would be appropriate to turn on physical scrolling using `scrollok()`.

**Soft Labels**

If desired, `curses` will manipulate the set of soft function-key labels that exist on many terminals. For those terminals that do not have soft labels, if you want to simulate them, `curses` will take over the bottom line of `stdscr`, reducing the size of `stdscr` and the variable `LINES`. `curses` standardizes on 8 labels of 8 characters each.

- slk\_init(*labfmt*)** In order to use soft labels, this routine must be called before `initscr()` or `newterm()` is called. If `initscr()` winds up using a line from `stdscr` to emulate the soft labels, then *labfmt* determines how the labels are arranged on the screen. Setting *labfmt* to 0 indicates that the labels are to be arranged in a 3-2-3 arrangement; 1 asks for a 4-4 arrangement.

- slk\_set(*labnum*, *label*, *labfmt*)** *labnum* is the label number, from 1 to 8. *label* is the string to be put on the label, up to 8 characters in length. A null string or a NULL pointer will put up a blank label. *labfmt* is one of 0, 1 or 2, to indicate whether the label is to be left-justified, centered, or right-justified within the label.

**slk\_refresh()**

- slk\_noutrefresh()** These routines correspond to the routines `wrefresh()` and `wnoutrefresh()`. Most applications would use `slk_noutrefresh()` because a `wrefresh()` will most likely soon follow.

- char \*slk\_label(*labnum*)** The current label for label number *labnum*, with leading and trailing blanks stripped, is returned.

**slk\_clear()** The soft labels are cleared from the screen.

**slk\_restore()** The soft labels are restored to the screen after a `slk_clear()`.

**slk\_touch()** All of the soft labels are forced to be output the next time a `slk_noutrefresh()` is performed.

**Low-Level curses Access**

The following routines give low-level access to various `curses` functionality. These routines typically would be used inside of library routines.

**def\_prog\_mode()**

- def\_shell\_mode()** Save the current terminal modes as the “program” (in `curses`) or “shell” (not in `curses`) state for use by the `reset_prog_mode()` and `reset_shell_mode()` routines. This is done automatically by `initscr()`.

**reset\_prog\_mode()**

- reset\_shell\_mode()** Restore the terminal to “program” (in `curses`) or “shell” (out of `curses`) state. These are done automatically by `endwin()` and `doupdate()` after an `endwin()`, so they normally would not be called.



- getsyx** (*y, x*) The current coordinates of the virtual screen cursor are returned in *y* and *x*. Like **getyx()**, the variables *y* and *x* do not take an **&** before them. If **leaveok()** is currently **TRUE**, then **-1, -1** will be returned. If lines may have been removed from the top of the screen using **ripoffline()** and the values are to be used beyond just passing them on to **setsyx()**, the value **y+stdscr->\_yoffset** should be used for those other uses.
- Note: **getsyx()** is a macro.
- setsyx** (*y, x*) The virtual screen cursor is set to *y, x*. If *y* and *x* are both **-1**, then **leaveok()** will be set. The two routines **getsyx()** and **setsyx()** are designed to be used by a library routine that manipulates **curses** windows but does not want to mess up the current position of the program's cursor. The library routine would call **getsyx()** at the beginning, do its manipulation of its own windows, do a **wnoutrefresh()** on its windows, call **setsyx()**, and then call **doupdate()**.
- ripoffline** (*line, init*) This routine provides access to the same facility that **slk\_init()** uses to reduce the size of the screen. **ripoffline()** must be called before **initscr()** or **newterm()** is called. If *line* is positive, a line will be removed from the top of **stdscr**; if negative, a line will be removed from the bottom. When this is done inside **initscr()**, the routine *init* is called with two arguments: a window pointer to the 1-line window that has been allocated and an integer with the number of columns in the window. Inside this initialization routine, the integer variables **LINES** and **COLS** (defined in **<curses.h>**) are not guaranteed to be accurate and **wrefresh()** or **doupdate()** must not be called. It is allowable to call **wnoutrefresh()** during the initialization routine.
- ripoffline()** can be called up to five times before calling **initscr()** or **newterm()**.
- scr\_dump** (*filename*) The current contents of the virtual screen are written to the file *filename*.
- scr\_restore** (*filename*) The virtual screen is set to the contents of *filename*, which must have been written using **scr\_dump()**. The next call to **doupdate()** will restore the screen to what it looked like in the dump file.
- scr\_init** (*filename*) The contents of *filename* are read in and used to initialize the **curses** data structures about what the terminal currently has on its screen. If the data is determined to be valid, **curses** will base its next update of the screen on this information rather than clearing the screen and starting from scratch. **scr\_init()** would be used after **initscr()** or a **system(3)** call to share the screen with another process that has done a **scr\_dump()** after its **endwin()** call. The data will be declared invalid if the time-stamp of the tty is old or the **terminfo(5V)** capability **nrrmc** is true.
- curs\_set** (*visibility*) The cursor is set to invisible, normal, or very visible for *visibility* equal to **0, 1** or **2**.
- draino** (*ms*) Wait until the output has drained enough that it will only take *ms* more milliseconds to drain completely.
- garbagedlines** (*win, begline, numlines*) This routine indicates to **curses** that a screen line is garbaged and should be thrown away before having anything written over the top of it. It could be used for programs such as editors that want a command to redraw just a single line. Such a command could be used in cases where there is a noisy communications line and redrawing the entire screen would be subject to even more communication noise. Just redrawing the single line gives some semblance of hope that it would show up unblemished. The current location of the window is used to determine which lines are to be redrawn.

**napms** (*ms*)            Sleep for *ms* milliseconds.

### Terminfo-Level Manipulations

These low-level routines must be called by programs that need to deal directly with the **terminfo(5V)** database to handle certain terminal capabilities, such as programming function keys. For all other functionality, **curses** routines are more suitable and their use is recommended.

Initially, **setupterm()** should be called. Note: **setupterm()** is automatically called by **initscr()** and **newterm()**. This will define the set of terminal-dependent variables defined in the **terminfo(5V)** database. The **terminfo(5V)** variables *lines* and *columns* (see **terminfo(5V)**) are initialized by **setupterm()** as follows: if the environment variables **LINES** and **COLUMNS** exist, their values are used. If the above environment variables do not exist, and the window sizes in rows and columns as returned by the **TIOCGWINSZ ioctl** are non-zero, those sizes are used. Otherwise, the values for *lines* and *columns* specified in the **terminfo(5V)** database are used.

The header files **<curses.h>** and **<term.h>** should be included, in this order, to get the definitions for these strings, numbers, and flags. Parameterized strings should be passed through **tparm()** to instantiate them. All **terminfo(5V)** strings (including the output of **tparm()**) should be printed with **tputs()** or **putp()**. Before exiting, **reset\_shell\_mode()** should be called to restore the tty modes. Programs that use cursor addressing should output **enter\_ca\_mode** upon startup and should output **exit\_ca\_mode** before exiting (see **terminfo(5V)**). Programs desiring shell escapes should call **reset\_shell\_mode()** and output **exit\_ca\_mode** before the shell is called and should output **enter\_ca\_mode** and call **reset\_prog\_mode()** after returning from the shell. Note: this is different from the **curses** routines (see **endwin()**).

**setupterm** (*term, fildes, errret*)

Reads in the **terminfo(5V)** database, initializing the **terminfo(5V)** structures, but does not set up the output virtualization structures used by **curses**. The terminal type is in the character string *term*; if *term* is NULL, the environment variable **TERM** will be used. All output is to the file descriptor *fildes*. If *errret* is not NULL, then **setupterm()** will return OK or ERR and store a status value in the integer pointed to by *errret*. A status of 1 in *errret* is normal, 0 means that the terminal could not be found, and -1 means that the **terminfo(5V)** database could not be found. If *errret* is NULL, **setupterm()** will print an error message upon finding an error and exit. Thus, the simplest call is '**setupterm** ((char \*)0, 1, (int \*)0)', which uses all the defaults.

The **terminfo(5V)** boolean, numeric and string variables are stored in a structure of type **TERMINAL**. After **setupterm()** returns successfully, the variable *cur\_term* (of type **TERMINAL \***) is initialized with all of the information that the **terminfo(5V)** boolean, numeric and string variables refer to. The pointer may be saved before calling **setupterm()** again. Further calls to **setupterm()** will allocate new space rather than reuse the space pointed to by *cur\_term*.

**set\_curterm** (*nterm*)    *nterm* is of type **TERMINAL \***. **set\_curterm()** sets the variable *cur\_term* to *nterm*, and makes all of the **terminfo(5V)** boolean, numeric and string variables use the values from *nterm*.

**del\_curterm** (*oterm*)    *oterm* is of type **TERMINAL \***. **del\_curterm()** frees the space pointed to by *oterm* and makes it available for further use. If *oterm* is the same as *cur\_term*, then references to any of the **terminfo(5V)** boolean, numeric and string variables thereafter may refer to invalid memory locations until another **setupterm()** has been called.

**restartterm** (*term, fildes, errret*)

Like **setupterm()** after a memory restore.

**char \*tparm** (*str, p<sub>1</sub>, p<sub>2</sub>, ..., p<sub>g</sub>*)

Instantiate the string *str* with parms *p<sub>i</sub>*. A pointer is returned to the result of *str* with the parameters applied.

- tputs** (*str, count, putc*) Apply padding to the string *str* and output it. *str* must be a **terminfo(5V)** string variable or the return value from **tparm()**, **tgetstr()**, **tigetstr()** or **tgoto()**. *count* is the number of lines affected, or 1 if not applicable. **putchar()** is a **putc(3s)**-like routine to which the characters are passed, one at a time.
- putp** (*str*) A routine that calls **tputs()** (*str*, 1, **putc(3s)**).
- vidputs** (*attrs, putc*) Output a string that puts the terminal in the video attribute mode *attrs*, which is any combination of the attributes listed below. The characters are passed to the **putc(3s)**-like routine **putc(3s)**.
- vidattr** (*attrs*) Like **vidputs()**, except that it outputs through **putc(3s)**.
- tigetflag** (*capname*) The value -1 is returned if *capname* is not a boolean capability.
- tigetnum** (*capname*) The value -2 is returned if *capname* is not a numeric capability.
- tigetstr** (*capname*) The value (char \*) -1 is returned if *capname* is not a string capability.

#### Termcap Emulation

These routines are included as a conversion aid for programs that use the **termcap(3X)** library. Their parameters are the same and the routines are emulated using the **terminfo(5V)** database.

- tgetent** (*bp, name*) Look up **termcap** entry for *name*. The emulation ignores the buffer pointer *bp*.
- tgetflag** (*codename*) Get the boolean entry for *codename*.
- tgetnum** (*codes*) Get numeric entry for *codename*.
- char \*tgetstr** (*codename, area*)  
Return the string entry for *codename*. If *area* is not NULL, then also store it in the buffer pointed to by *area* and advance *area*. **tputs()** should be used to output the returned string.
- char \*tgoto** (*cap, col, row*)  
Instantiate the parameters into the given capability. The output from this routine is to be passed to **tputs()**.
- tputs** (*str, affcnt, putc*) See **tputs()** above, under **Terminfo-Level Manipulations**.

#### Miscellaneous

- char \*keyname** (*c*) A character string corresponding to the key *c* is returned.
- filter()** This routine is one of the few that is to be called before **initscr()** or **newterm()** is called. It arranges things so that **curses** thinks that there is a 1-line screen. **curses** will not use any terminal capabilities that assume that they know what line on the screen the cursor is on.

#### Use of curscr

The special window **curscr** can be used in only a few routines. If the window argument to **clearok()** is **curscr**, the next call to **wrefresh()** with any window will cause the screen to be cleared and repainted from scratch. If the window argument to **wrefresh()** is **curscr**, the screen is immediately cleared and repainted from scratch. This is how most programs would implement a "repaint-screen" routine. The source window argument to **overlay()**, **overwrite()**, and **copywin** may be **curscr**, in which case the current contents of the virtual terminal screen will be accessed.

#### Obsolete Calls

Various routines are provided to maintain compatibility in programs written for older versions of the **curses** library. These routines are all emulated as indicated below.

- crmode()** Replaced by **cbreak()**.
- fixterm()** Replaced by **reset\_prog\_mode()**.
- nocrmode()** Replaced by **nocbreak()**.

**resetterm()** Replaced by **reset\_shell\_mode()**.  
**saveterm()** Replaced by **def\_prog\_mode()**.  
**setterm()** Replaced by **setupterm()**.

**SYSTEM V ATTRIBUTES**

The following video attributes, defined in `< curses.h >`, can be passed to the routines **attron()**, **attroff()**, and **attrset()**, or OR'ed with the characters passed to **addch()**.

**A\_STANDOUT** Terminal's best highlighting mode  
**A\_UNDERLINE** Underlining  
**A\_REVERSE** Reverse video  
**A\_BLINK** Blinking  
**A\_DIM** Half bright  
**A\_BOLD** Extra bright or bold  
**A\_ALTCHARSET** Alternate character set

**A\_CHARTEXT** Bit-mask to extract character (described under **winch**)  
**A\_ATTRIBUTES** Bit-mask to extract attributes (described under **winch**)  
**A\_NORMAL** Bit mask to reset all attributes off  
(for example: `'attrset (A_NORMAL)'`)

**SYSTEM V FUNCTION KEYS**

The following function keys, defined in `< curses.h >`, might be returned by **getch()** if **keypad()** has been enabled. Note: not all of these may be supported on a particular terminal if the terminal does not transmit a unique code when the key is pressed or the definition for the key is not present in the **terminfo(5V)** database.

<i>Name</i>	<i>Value</i>	<i>Key name</i>
<b>KEY_BREAK</b>	0401	break key (unreliable)
<b>KEY_DOWN</b>	0402	The four arrow keys ...
<b>KEY_UP</b>	0403	
<b>KEY_LEFT</b>	0404	
<b>KEY_RIGHT</b>	0405	...
<b>KEY_HOME</b>	0406	Home key (upward+left arrow)
<b>KEY_BACKSPACE</b>	0407	backspace (unreliable)
<b>KEY_F0</b>	0410	Function keys. Space for 64 keys is reserved.
<b>KEY_F(n)</b>	( <b>KEY_F0</b> +(n))	Formula for $f_n$ .
<b>KEY_DL</b>	0510	Delete line
<b>KEY_IL</b>	0511	Insert line
<b>KEY_DC</b>	0512	Delete character
<b>KEY_IC</b>	0513	Insert char or enter insert mode
<b>KEY_EIC</b>	0514	Exit insert char mode
<b>KEY_CLEAR</b>	0515	Clear screen
<b>KEY_EOS</b>	0516	Clear to end of screen
<b>KEY_EOL</b>	0517	Clear to end of line
<b>KEY_SF</b>	0520	Scroll 1 line forward
<b>KEY_SR</b>	0521	Scroll 1 line backwards (reverse)
<b>KEY_NPAGE</b>	0522	Next page
<b>KEY_PPAGE</b>	0523	Previous page
<b>KEY_STAB</b>	0524	Set TAB
<b>KEY_CTAB</b>	0525	Clear TAB
<b>KEY_CATAB</b>	0526	Clear all TAB characters
<b>KEY_ENTER</b>	0527	Enter or send
<b>KEY_SRESET</b>	0530	soft (partial) reset

KEY_RESET	0531	reset or hard reset
KEY_PRINT	0532	print or copy
KEY_LL	0533	home down or bottom (lower left) keypad is arranged like this: A1 up A3 left B2 right C1 down C3
KEY_A1	0534	Upper left of keypad
KEY_A3	0535	Upper right of keypad
KEY_B2	0536	Center of keypad
KEY_C1	0537	Lower left of keypad
KEY_C3	0540	Lower right of keypad
KEY_BTAB	0541	Back TAB key
KEY_BEG	0542	beg(inning) key
KEY_CANCEL	0543	cancel key
KEY_CLOSE	0544	close key
KEY_COMMAND	0545	cmd (command) key
KEY_COPY	0546	copy key
KEY_CREATE	0547	create key
KEY_END	0550	end key
KEY_EXIT	0551	exit key
KEY_FIND	0552	find key
KEY_HELP	0553	help key
KEY_MARK	0554	mark key
KEY_MESSAGE	0555	message key
KEY_MOVE	0556	move key
KEY_NEXT	0557	next object key
KEY_OPEN	0560	open key
KEY_OPTIONS	0561	options key
KEY_PREVIOUS	0562	previous object key
KEY_REDO	0563	redo key
KEY_REFERENCE	0564	ref(erence) key
KEY_REFRESH	0565	refresh key
KEY_REPLACE	0566	replace key
KEY_RESTART	0567	restart key
KEY_RESUME	0570	resume key
KEY_SAVE	0571	save key
KEY_SBEG	0572	shifted beginning key
KEY_SCANCEL	0573	shifted cancel key
KEY_SCOMMAND	0574	shifted command key
KEY_SCOPY	0575	shifted copy key
KEY_SCREATE	0576	shifted create key
KEY_SDC	0577	shifted delete char key
KEY_SDL	0600	shifted delete line key
KEY_SELECT	0601	select key
KEY_SEND	0602	shifted end key
KEY_SEOL	0603	shifted clear line key
KEY_SEXIT	0604	shifted exit key
KEY_SFIND	0605	shifted find key
KEY_SHELP	0606	shifted help key
KEY_SHOME	0607	shifted home key
KEY_SIC	0610	shifted input key
KEY_SLEFT	0611	shifted left arrow key

KEY_SMESSAGE	0612	shifted message key
KEY_SMOVE	0613	shifted move key
KEY_SNEXT	0614	shifted next key
KEY_SOPTIONS	0615	shifted options key
KEY_SPREVIOUS	0616	shifted prev key
KEY_SPRINT	0617	shifted print key
KEY_SREDO	0620	shifted redo key
KEY_SREPLACE	0621	shifted replace key
KEY_SRIGHT	0622	shifted right arrow
KEY_SRSUME	0623	shifted resume key
KEY_SSAVE	0624	shifted save key
KEY_SSUSPEND	0625	shifted suspend key
KEY_SUNDO	0626	shifted undo key
KEY_SUSPEND	0627	suspend key
KEY_UNDO	0630	undo key

**LINE GRAPHICS**

The following variables may be used to add line-drawing characters to the screen with `waddch`. When defined for the terminal, the variable will have the `A_ALTCHARSET` bit turned on. Otherwise, the default character listed below will be stored in the variable. The names were chosen to be consistent with the DEC VT100 nomenclature.

<i>Name</i>	<i>Default</i>	<i>Glyph Description</i>
ACS_ULCORNER	+	upper left corner
ACS_LLCORNER	+	lower left corner
ACS_URCORNER	+	upper right corner
ACS_LRCORNER	+	lower right corner
ACS_RTEE	+	right tee (┌)
ACS_LTEE	+	left tee (└)
ACS_BTEE	+	bottom tee (┘)
ACS_TTEE	+	top tee (┐)
ACS_HLINE	-	horizontal line
ACS_VLINE		vertical line
ACS_PLUS	+	plus
ACS_S1	-	scan line 1
ACS_S9	-	scan line 9
ACS_DIAMOND	+	diamond
ACS_CKBOARD	:	checker board (stipple)
ACS_DEGREE	'	degree symbol
ACS_PLMINUS	#	plus/minus
ACS_BULLET	o	bullet
ACS_LARROW	<	arrow pointing left
ACS_RARROW	>	arrow pointing right
ACS_DARROW	v	arrow pointing down
ACS_UARROW	^	arrow pointing up
ACS_BOARD	#	board of squares
ACS_LANTERN	#	lantern symbol
ACS_BLOCK	#	solid square block

**RETURN VALUES**

Unless otherwise noted in the preceding routine descriptions, all routines return:

OK     on success.  
ERR    on failure.

**SYSTEM V RETURN VALUES**

All macros return the value of their **w** version, except **setscrreg()**, **wsetscrreg()**, **getsyx()**, **getyx()**, **getbegy()**, **getmaxyx()**, which return no useful value.

Routines that return pointers always return (*type \**) NULL on failure.

**FILES**

**.login**

**.profile**

**SYSTEM V FILES**

**/usr/share/lib/terminfo**

**SEE ALSO**

**cc(1V)**, **ld(1)**, **ioctl(2)**, **getenv(3V)**, **plot(3X)**, **printf(3V)**, **putc(3S)**, **scanf(3V)**, **stdio(3V)**, **system(3)**, **varargs(3)**, **vprintf(3V)**, **termio(4)**, **tty(4)**, **term(5V)**, **termcap(5)**, **terminfo(5V)**, **tic(8V)**

**SYSTEM V WARNINGS**

The plotting library **plot(3X)** and the curses library **curses(3V)** both use the names **erase()** and **move()**. The **curses** versions are macros. If you need both libraries, put the **plot(3X)** code in a different source file than the **curses(3V)** code, and/or **'#undef move'** and **'#undef erase'** in the **plot(3X)** code.

Between the time a call to **initscr()** and **endwin()** has been issued, use only the routines in the **curses** library to generate output. Using system calls or the "standard I/O package" (see **stdio(3V)**) for output during that time can cause unpredictable results.

**NAME**

`cuserid` – get character login name of the user

**SYNOPSIS**

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

```
char *cuserid(s)
```

```
char *s;
```

**DESCRIPTION**

`cuserid()` returns a pointer to a string representing the login name under which the owner of the current process is logged in. If *s* is a NULL pointer, this string is placed in an internal static area, the address of which is returned. Otherwise, *s* is assumed to point to an array of at least `L_cuserid` characters; the representation is left in this array. The constant `L_cuserid` is defined in the `<stdio.h>` header file.

**SEE ALSO**

`cc(1V)`, `ld(1)`, `getlogin(3V)`, `getpwent(3V)`

**RETURN VALUES**

`cuserid()` returns a pointer to the login name on success. On failure, `cuserid()` returns NULL, and if *s* is not NULL, places a null character ('\0') at `s[0]`.

**NOTES**

The internal static area to which `cuserid()` writes when *s* is NULL will be overwritten by a subsequent call to `getpwnam()` (see `getpwent(3V)`).

A compatibility problem has been identified with the `cuserid()` function. The traditional version of this library routine in SunOS Release 3.2 and later releases and all System V releases calls the `getlogin()` function, and if it fails uses the `getpwuid()` function to try to return a name associated with the real user ID associated with the calling process. POSIX.1 requires that the `cuserid()` function try to return a name associated with the effective user ID associated with the calling process. Although this usually yields the same results, use of set-uid programs may yield different results.

A binding interpretation has been issued by IEEE saying that the POSIX.1 functionality has to be provided for compliance with POSIX.1. However, balloting on the first update to POSIX.1, P1003.1a, has led to the removal of the `cuserid()` function from the standard. (This is the state in the second recirculation ballot of P1003.1a dated 11 December 1989.) The objections leading to this resolution had both users and implementors arguing for the historical version and for the version specified by POSIX.1. The only way to reach consensus appears to be to remove the function from the standard.

To further complicate the issue, System V Release 4.0 has kept the traditional version of `cuserid()`. XPG3 specifies the POSIX.1 version of `cuserid()`, but the test suite for conformance to XPG3 promises to accept either implementation. Both of these are anticipating the final approval of P1003.1a as a standard with the `cuserid()` function removed. Since we also expect the `cuserid()` function to be dropped from the standard when P1003.1a is approved, SunOS Release 4.1 provides the traditional `cuserid()` function in the C library. However, for users that need the version specified by POSIX.1, it is provided in a POSIX library available in the System V environment. This library can be accessed by specifying `-lposix` on the `cc(1V)` or `ld(1)` command line.



## NAME

dbm, dbmopen, dbmclose, fetch, store, delete, firstkey, nextkey – data base subroutines

## SYNOPSIS

```
#include <dbm.h>

typedef struct {
    char *dptr;
    int dsize;
} datum;

dbmopen(file)
char *file;

dbmclose()

datum fetch(key)
datum key;

store(key, content)
datum key, content;

delete(key)
datum key;

datum firstkey()

datum nextkey(key)
datum key;
```

## DESCRIPTION

Note: the **dbm()** library has been superseded by **ndbm(3)**, and is now implemented using **ndbm()**.

These functions maintain key/content pairs in a data base. The functions will handle very large (a billion blocks) databases and will access a keyed item in one or two file system accesses. The functions are obtained with the loader option **-ldb**.

*keys* and *contents* are described by the **datum** typedef. A **datum** specifies a string of *dsize* bytes pointed to by *dptr*. Arbitrary binary data, as well as normal ASCII strings, are allowed. The data base is stored in two files. One file is a directory containing a bit map and has **.dir** as its suffix. The second file contains all data and has **.pag** as its suffix.

Before a database can be accessed, it must be opened by **dbmopen**. At the time of this call, the files **file.dir** and **file.pag** must exist. (An empty database is created by creating zero-length **.dir** and **.pag** files.)

A database may be closed by calling **dbmclose**. You must close a database before opening a new one.

Once open, the data stored under a key is accessed by **fetch()** and data is placed under a key by **store**. A key (and its associated contents) is deleted by **delete**. A linear pass through all keys in a database may be made, in an (apparently) random order, by use of **firstkey()** and **nextkey**. **firstkey()** will return the first key in the database. With any key **nextkey()** will return the next key in the database. This code will traverse the data base:

```
for (key = firstkey(); key.dptr != NULL; key = nextkey(key))
```

## SEE ALSO

**ar(1V)**, **cat(1V)**, **cp(1)**, **tar(1)**, **ndbm(3)**

## DIAGNOSTICS

All functions that return an **int** indicate errors with negative values. A zero return indicates no error. Routines that return a **datum** indicate errors with a NULL (0) *dptr*.

**BUGS**

The **.pag** file will contain holes so that its apparent size is about four times its actual content. Older versions of the UNIX operating system may create real file blocks for these holes when touched. These files cannot be copied by normal means (**cp(1)**, **cat(1V)**, **tar(1)**, **ar(1V)**) without filling in the holes.

*dptr* pointers returned by these subroutines point into static storage that is changed by subsequent calls.

The sum of the sizes of a key/content pair must not exceed the internal block size (currently 1024 bytes). Moreover all key/content pairs that hash together must fit on a single block. **store()** will return an error in the event that a disk block fills with inseparable data.

**delete()** does not physically reclaim file space, although it does make it available for reuse.

The order of keys presented by **firstkey()** and **nextkey()** depends on a hashing function, not on anything interesting.

There are no interlocks and no reliable cache flushing; thus concurrent updating and reading is risky.

## NAME

`decimal_to_single`, `decimal_to_double`, `decimal_to_extended` – convert decimal record to floating-point value

## SYNOPSIS

```
#include <floatingpoint.h>

void decimal_to_single(px, pm, pd, ps)
single *px ;
decimal_mode *pm;
decimal_record *pd;
fp_exception_field_type *ps;

void decimal_to_double(px, pm, pd, ps)
double *px ;
decimal_mode *pm;
decimal_record *pd;
fp_exception_field_type *ps;

void decimal_to_extended(px, pm, pd, ps)
extended *px ;
decimal_mode *pm;
decimal_record *pd;
fp_exception_field_type *ps;
```

## DESCRIPTION

The `decimal_to_floating()` functions convert the decimal record at `*pd` into a floating-point value at `*px`, observing the modes specified in `*pm` and setting exceptions in `*ps`. If there are no IEEE exceptions, `*ps` will be zero.

`pd->sign` and `pd->fpclass` are always taken into account. `pd->exponent` and `pd->ds` are used when `pd->fpclass` is `fp_normal` or `fp_subnormal`. In these cases `pd->ds` must contain one or more ascii digits followed by a null character. `*px` is set to a correctly rounded approximation to

$$(\text{pd->sign}) * (\text{pd->ds}) * 10^{(\text{pd->exponent})}$$

Thus if `pd->exponent == -2` and `pd->ds == "1234"`, `*px` will get 12.34 rounded to storage precision. `pd->ds` cannot have more than `DECIMAL_STRING_LENGTH-1` significant digits because one character is used to terminate the string with a null character. If `pd->more != 0` on input then additional nonzero digits follow those in `pd->ds`; `fp_inexact` is set accordingly on output in `*ps`.

`*px` is correctly rounded according to the IEEE rounding modes in `pm->rd`. `*ps` is set to contain `fp_inexact`, `fp_underflow`, or `fp_overflow` if any of these arise.

`pd->ndigits`, `pm->df`, and `pm->ndigits` are not used.

`strtod(3)`, `scanf(3V)`, `fscanf()`, and `sscanf()` all use `decimal_to_double()`.

## SEE ALSO

`scanf(3V)`, `strtod(3)`

## NAME

`des_crypt`, `ecb_crypt`, `cbc_crypt`, `des_setparity` – fast DES encryption

## SYNOPSIS

```
#include <des_crypt.h>

int ecb_crypt(key, data, datalen, mode)
char *key;
char *data;
unsigned datalen;
unsigned mode;

int cbc_crypt(key, data, datalen, mode, ivec)
char *key;
char *data;
unsigned datalen;
unsigned mode;
char *ivec;

void des_setparity(key)
char *key;
```

## DESCRIPTION

`ecb_crypt()` and `cbc_crypt()` implement the NBS DES (Data Encryption Standard). These routines are faster and more general purpose than `crypt(3)`. They also are able to utilize DES hardware if it is available. `ecb_crypt()` encrypts in ECB (Electronic Code Book) mode, which encrypts blocks of data independently. `cbc_crypt()` encrypts in CBC (Cipher Block Chaining) mode, which chains together successive blocks. CBC mode protects against insertions, deletions and substitutions of blocks. Also, regularities in the clear text will not appear in the cipher text.

Here is how to use these routines. The first parameter, *key*, is the 8-byte encryption key with parity. To set the key's parity, which for DES is in the low bit of each byte, use `des_setparity`. The second parameter, *data*, contains the data to be encrypted or decrypted. The third parameter, *datalen*, is the length in bytes of *data*, which must be a multiple of 8. The fourth parameter, *mode*, is formed by OR'ing together some things. For the encryption direction 'or' in either `DES_ENCRYPT` or `DES_DECRYPT`. For software versus hardware encryption, 'or' in either `DES_HW` or `DES_SW`. If `DES_HW` is specified, and there is no hardware, then the encryption is performed in software and the routine returns `DESERR_NOHWDEVICE`. For `cbc_crypt`, the parameter *ivec* is the 8-byte initialization vector for the chaining. It is updated to the next initialization vector upon return.

## SEE ALSO

`des(1)`, `crypt(3)`

## DIAGNOSTICS

<code>DESERR_NONE</code>	No error.
<code>DESERR_NOHWDEVICE</code>	Encryption succeeded, but done in software instead of the requested hardware.
<code>DESERR_HWERR</code>	An error occurred in the hardware or driver.
<code>DESERR_BADPARAM</code>	Bad parameter to routine.

Given a result status *stat*, the macro `DES_FAILED(stat)` is false only for the first two statuses.

## RESTRICTIONS

These routines are not available for export outside the U.S.

## NAME

directory, opendir, readdir, telldir, seekdir, rewinddir, closedir – directory operations

## SYNOPSIS

```
#include <dirent.h>

DIR *opendir(dirname)
char *dirname;

struct dirent *readdir(dirp)
DIR *dirp;

long telldir(dirp)
DIR *dirp;

void seekdir(dirp, loc)
DIR *dirp;
long loc;

void rewinddir(dirp)
DIR *dirp;

int closedir(dirp)
DIR *dirp;
```

## SYSTEM V SYNOPSIS

For XPG2 conformance, use:

```
#include <sys/dirent.h>
```

## DESCRIPTION

**opendir()** opens the directory named by *dirname* and associates a *directory stream* with it. **opendir()** returns a pointer to be used to identify the directory stream in subsequent operations. A NULL pointer is returned if *dirname* cannot be accessed or is not a directory, or if it cannot **malloc(3V)** enough memory to hold the whole thing.

**readdir()** returns a pointer to the next directory entry. It returns NULL upon reaching the end of the directory or detecting an invalid **seekdir()** operation.

**telldir()** returns the current location associated with the named directory stream.

**seekdir()** sets the position of the next **readdir()** operation on the directory stream. The new position reverts to the one associated with the directory stream when the **telldir()** operation was performed. Values returned by **telldir()** are good only for the lifetime of the DIR pointer from which they are derived. If the directory is closed and then reopened, the **telldir()** value may be invalidated due to undetected directory compaction. It is safe to use a previous **telldir()** value immediately after a call to **opendir()** and before any calls to **readdir()**.

**rewinddir()** resets the position of the named directory stream to the beginning of the directory. It also causes the directory stream to refer to the current state of the corresponding directory, as a call to **opendir()** would have done.

**closedir()** closes the named directory stream and frees the structure associated with the DIR pointer.

**RETURN VALUES**

**opendir()** returns a pointer to an object of type **DIR** on success. On failure, it returns **NULL** and sets **errno** to indicate the error.

**readdir()** returns a pointer to an object of type **struct dirent** on success. On failure, it returns **NULL** and sets **errno** to indicate the error. When the end of the directory is encountered, **readdir()** returns **NULL** and leaves **errno** unchanged.

**closedir()** returns:

0 on success.

-1 on failure and sets **errno** to indicate the error.

**telldir()** returns the current location associated with the specified directory stream.

**ERRORS**

If any of the following conditions occur, **opendir()** sets **errno** to:

**EACCES** Search permission is denied for a component of *dirname*.  
Read permission is denied for *dirname*.

**ENAMETOOLONG** The length of *dirname* exceeds {**PATH\_MAX**}.  
A pathname component is longer than {**NAME\_MAX**} (see **sysconf(2V)**) while {**\_POSIX\_NO\_TRUNC**} is in effect (see **pathconf(2V)**).

**ENOENT** The named directory does not exist.

**ENOTDIR** A component of *dirname* is not a directory.

for each of the following conditions, when the condition is detected, **opendir()** sets **errno** to one of the following:

**EMFILE** Too many file descriptors are currently open for the process.

**ENFILE** Too many file descriptors are currently open in the system.

For each of the following conditions, when the condition is detected, **readdir()** sets **errno** to the following:

**EBADF** *dirp* does not refer to an open directory stream.

For each of the following conditions, when the condition is detected, **closedir()** sets **errno** to the following:

**EBADF** *dirp* does not refer to an open directory stream.

**SYSTEM V ERRORS**

In addition to the above, **opendir()** may set **errno** to the following:

**ENOENT** *dirname* points to an empty string.

**EXAMPLES**

Sample code which searches a directory for entry "name" is:

```

dirp = opendir(".");
for (dp = readdir(dirp); dp != NULL; dp = readdir(dirp))
    if (!strcmp(dp->d_name, name)) {
        closedir (dirp);
        return FOUND;
    }
closedir (dirp);
return NOT_FOUND;

```

**SEE ALSO**

**close(2V), lseek(2V), open(2V), read(2V), getwd(3), malloc(3V), dir(5)**

**NOTES**

The **directory** library routines now use a new include file, **<dirent.h>**. This replaces the file, **<sys/dir.h>**, used in previous releases. Furthermore, with the use of this new file, the **readdir()** routine returns directory entries whose structure is named **struct dirent** rather than **struct direct** as before. The file **<sys/dir.h>** is retained in the current SunOS release for purposes of backwards source code compatibility; programs which use the **directory()** library and **<sys/dir.h>** will continue to compile and run without source code modifications. However, existing programs should convert to the use of the new include file, **<dirent.h>**, as **<sys/dir.h>** will be removed in a future major release.

The *X/Open Portability Guide, issue 2* (XPG2) requires **<sys/dirent.h>** rather than **<dirent.h>**. **/usr/xpg2include/sys/dirent.h** is functionally equivalent to **/usr/include/dirent.h**. In future SunOS releases, X/Open conformance will require **<dirent.h>**.

## NAME

`dlopen`, `dlsym`, `dlerror`, `dlclose` – simple programmatic interface to the dynamic linker

## SYNOPSIS

```
#include <dlfcn.h>

void *dlopen(path, mode)
char *path; int mode;

void *dlsym(handle, symbol)
void *handle; char *symbol;

char *dlerror()

int dlclose(handle);
void *handle;
```

## DESCRIPTION

These functions provide a simple programmatic interface to the services of the dynamic link-editor. Operations are provided to add a new shared object to an program's address space, obtain the address bindings of symbols defined by such objects, and to remove such objects when their use is no longer required.

`dlopen()` provides access to the shared object in *path*, returning a descriptor that can be used for later references to the object in calls to `dlsym()` and `dlclose()`. If *path* was not in the address space prior to the call to `dlopen()`, then it will be placed in the address space, and if it defines a function with the name `_init` that function will be called by `dlopen()`. If, however, *path* has already been placed in the address space in a previous call to `dlopen()`, then it will not be added a second time, although a count of `dlopen()` operations on *path* will be maintained. *mode* is an integer containing flags describing options to be applied to the opening and loading process — it is reserved for future expansion and must always have the value 1. A null pointer supplied for *path* is interpreted as a reference to the “main” executable of the process. If `dlopen()` fails, it will return a null pointer.

`dlsym()` returns the address binding of the symbol described in the null-terminated character string *symbol* as it occurs in the shared object identified by *handle*. The symbols exported by objects added to the address space by `dlopen()` can be accessed *only* through calls to `dlsym()`, such symbols do not supersede any definition of those symbols already present in the address space when the object is loaded, nor are they available to satisfy “normal” dynamic linking references. `dlsym()` returns a null pointer if the symbol can not be found. A null pointer supplied as the value of *handle* is interpreted as a reference to the executable from which the call to `dlsym()` is being made — thus a shared object can reference its own symbols.

`dlerror` returns a null-terminated character string describing the last error that occurred during a `dlopen()`, `dlsym()`, or `dlclose()`. If no such error has occurred, then `dlerror()` will return a null pointer. At each call to `dlerror()`, the “last error” indication will be reset, thus in the case of two calls to `dlerror()`, and where the second call follows the first immediately, the second call will always return a null pointer.

`dlclose()` deletes a reference to the shared object referenced by *handle*. If the reference count drops to 0, then if the object referenced by *handle* defines a function `_fini`, that function will be called, the object removed from the address space, and *handle* destroyed. If `dlclose()` is successful, it will return a value of 0. A failing call to `dlclose()` will return a non-zero value.

The object-intrinsic functions `_init` and `_fini` are called with no arguments and treated as though their types were `void`.

These functions are obtained by specifying `-ldl` as an option to `ld(1)`.

## SEE ALSO

`ld(1)`, `link(5)`



## NAME

**drand48**, **erand48**, **lrand48**, **rand48**, **mrnd48**, **jrand48**, **srand48**, **seed48**, **lcong48** – generate uniformly distributed pseudo-random numbers

## SYNOPSIS

```
double drand48()

double erand48(xsubi)
unsigned short xsubi[3];

long lrand48()

long rand48(xsubi)
unsigned short xsubi[3];

long mrnd48()

long jrand48(xsubi)
unsigned short xsubi[3];

void srand48(seedval)
long seedval;

unsigned short *seed48(seed16v)
unsigned short seed16v[3];

void lcong48(param)
unsigned short param[7];
```

## DESCRIPTION

This family of functions generates pseudo-random numbers using the well-known linear congruential algorithm and 48-bit integer arithmetic.

**drand48()** and **erand48()** return non-negative double-precision floating-point values uniformly distributed over the interval [0.0, 1.0).

**lrand48()** and **rand48()** return non-negative long integers uniformly distributed over the interval [0, 2<sup>31</sup>).

**mrnd48()** and **jrand48()** return signed long integers uniformly distributed over the interval [-2<sup>31</sup>, 2<sup>31</sup>).

**srand48()**, **seed48()**, and **lcong48()** are initialization entry points, one of which should be invoked before either **drand48()**, **lrand48()**, or **mrnd48()** is called. Although it is not recommended practice, constant default initializer values will be supplied automatically if **drand48()**, **lrand48()**, or **mrnd48()** is called without a prior call to an initialization entry point. **erand48()**, **rand48()**, and **jrand48()** do not require an initialization entry point to be called first.

All the routines work by generating a sequence of 48-bit integer values,  $X_i$ , according to the linear congruential formula

$$X_{n+1} = (aX_n + c)_{\text{mod } m} \quad n \geq 0.$$

The parameter  $m = 2^{48}$ ; hence 48-bit integer arithmetic is performed. Unless **lcong48()** has been invoked, the multiplier value  $a$  and the addend value  $c$  are given by

$$a = 5DEECE66D_{16} = 273673163155_8$$

$$c = B_{16} = 13_8.$$

The value returned by any of the functions **drand48()**, **erand48()**, **lrand48()**, **rand48()**, **mrnd48()**, or **jrand48()** is computed by first generating the next 48-bit  $X_i$  in the sequence. Then the appropriate number of bits, according to the type of data item to be returned, are copied from the high-order (leftmost) bits of  $X_i$  and transformed into the returned value.

**drand48()**, **lrand48()**, and **mrnd48()** store the last 48-bit  $X_i$  generated in an internal buffer; that is why they must be initialized prior to being invoked. The functions **erand48()**, **rand48()**, and **jrand48()** require the calling program to provide storage for the successive  $X_i$  values in the array specified as an

argument when the functions are invoked. That is why these routines do not have to be initialized; the calling program merely has to place the desired initial value of  $X_i$  into the array and pass it as an argument. By using different arguments, functions `erand48()`, `rand48()`, and `jrand48()` allow separate modules of a large program to generate several *independent* streams of pseudo-random numbers, that is, the sequence of numbers in each stream will *not* depend upon how many times the routines have been called to generate numbers for the other streams.

The initializer function `srand48()` sets the high-order 32 bits of  $X_i$  to the 32 bits contained in its argument. The low-order 16 bits of  $X_i$  are set to the arbitrary value  $330E_{16}$ .

The initializer function `seed48()` sets the value of  $X_i$  to the 48-bit value specified in the argument array. In addition, the previous value of  $X_i$  is copied into a 48-bit internal buffer, used only by `seed48()`, and a pointer to this buffer is the value returned by `seed48()`. This returned pointer, which can just be ignored if not needed, is useful if a program is to be restarted from a given point at some future time — use the pointer to get at and store the last  $X_i$  value, and then use this value to reinitialize via `seed48()` when the program is restarted.

The initialization function `lcg48()` allows the user to specify the initial  $X_i$ , the multiplier value  $a$ , and the addend value  $c$ . Argument array elements `param[0-2]` specify  $X_i$ , `param[3-5]` specify the multiplier  $a$ , and `param[6]` specifies the 16-bit addend  $c$ . After `lcg48()` has been called, a subsequent call to either `srand48()` or `seed48()` will restore the “standard” multiplier and addend values,  $a$  and  $c$ , specified on the previous page.

**SEE ALSO**

`rand(3V)`

## NAME

econvert, fconvert, gconvert, seconvert, sfconvert, sgconvert, ecvt, fcvt, gcvt – output conversion

## SYNOPSIS

```
#include <floatingpoint.h>
```

```
char *econvert(value, ndigit, decpt, sign, buf)
double value;
int ndigit, *decpt, *sign;
char *buf;
```

```
char *fconvert(value, ndigit, decpt, sign, buf)
double value;
int ndigit, *decpt, *sign;
char *buf;
```

```
char *gconvert(value, ndigit, trailing, buf)
double value;
int ndigit;
int trailing;
char *buf;
```

```
char *seconvert(value, ndigit, decpt, sign, buf)
single *value;
int ndigit, *decpt, *sign;
char *buf;
```

```
char *sfconvert(value, ndigit, decpt, sign, buf)
single *value;
int ndigit, *decpt, *sign;
char *buf;
```

```
char *sgconvert(value, ndigit, trailing, buf)
single *value;
int ndigit;
int trailing;
char *buf;
```

```
char *ecvt(value, ndigit, decpt, sign)
double value;
int ndigit, *decpt, *sign;
```

```
char *fcvt(value, ndigit, decpt, sign)
double value;
int ndigit, *decpt, *sign;
```

```
char *gcvt(value, ndigit, buf)
double value;
int ndigit;
char *buf;
```

## DESCRIPTION

`econvert()` converts the *value* to a null-terminated string of *ndigit* ASCII digits in *buf* and returns a pointer to *buf*. *buf* should contain at least *ndigit+1* characters. The position of the radix character relative to the beginning of the string is stored indirectly through *decpt*. Thus *buf* == "314" and *\*decpt* == 1 corresponds to the numerical value 3.14, while *buf* == "314" and *\*decpt* == -1 corresponds to the numerical value .0314. If the sign of the result is negative, the word pointed to by *sign* is nonzero; otherwise it is zero. The least significant digit is rounded.

**fconvert** works much like **econvert**, except that the correct digit has been rounded as if for **sprintf(%w.nf)** output with  $n=ndigit$  digits to the right of the radix character. *ndigit* can be negative to indicate rounding to the left of the radix character. The return value is a pointer to *buf*. *buf* should contain at least  $310+max(0,ndigit)$  characters to accommodate any double-precision *value*.

**gconvert()** converts the *value* to a null-terminated ASCII string in *buf* and returns a pointer to *buf*. It produces *ndigit* significant digits in fixed-decimal format, like **sprintf(%w.nf)**, if possible, and otherwise in floating-decimal format, like **sprintf(%w.ne)**; in either case *buf* is ready for printing, with sign and exponent. The result corresponds to that obtained by

```
(void) sprintf(buf, "%w.ng", value);
```

If *trailing*=0, trailing zeros and a trailing point are suppressed, as in **sprintf(%g)**. If *trailing*!=0, trailing zeros and a trailing point are retained, as in **sprintf(%#g)**.

**seconvert**, **sfconvert**, and **sgconvert()** are single-precision versions of these functions, and are more efficient than the corresponding double-precision versions. A pointer rather than the value itself is passed to avoid C's usual conversion of single-precision arguments to double.

**ecvt()** and **fcvt()** are obsolete versions of **econvert()** and **fconvert()** that create a string in a static data area, overwritten by each call, and return values that point to that static data. These functions are therefore not reentrant.

**gcvt()** is an obsolete version of **gconvert()** that always suppresses trailing zeros and point.

IEEE Infinities and NaNs are treated similarly by these functions. "NaN" is returned for NaN, and "Inf" or "Infinity" for Infinity. The longer form is produced when *ndigit* >= 8.

The radix character is determined by the current setting of the program's locale (category LC\_NUMERIC). In the "C" locale or if the locale is undefined, the radix character defaults to a period '.'.

**SEE ALSO**

**printf(3V)**

**NAME**

*end*, *etext*, *edata* – last locations in program

**SYNOPSIS**

**extern *end*;**  
**extern *etext*;**  
**extern *edata*;**

**DESCRIPTION**

These names refer neither to routines nor to locations with interesting contents. The address of *etext* is the first address above the program text, *edata* above the initialized data region, and *end*() above the uninitialized data region.

When execution begins, the program break (the first location beyond the data) coincides with *end*, but it is reset by the routines *brk*(2), *malloc*(3V), standard input/output (*stdio*(3V)), the profile (*-p*) option of *cc*(1V), and so on. Thus, the current value of the program break should be determined by *sbrk*(0) (see *brk*(2)).

**SEE ALSO**

*cc*(1V), *brk*(2), *malloc*(3V), *stdio*(3V)

**NAME**

ethers, ether\_ntoa, ether\_aton, ether\_ntohost, ether\_hostton, ether\_line – Ethernet address mapping operations

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/socket.h>
#include <net/if.h>
#include <netinet/in.h>
#include <netinet/if_ether.h>

char *
ether_ntoa(e)
struct ether_addr *e;

struct ether_addr *ether_aton(s)
char *s;

ether_ntohost(hostname, e)
char *hostname;
struct ether_addr *e;

ether_hostton(hostname, e)
char *hostname;
struct ether_addr *e;

ether_line(l, e, hostname)
char *l;
struct ether_addr *e;
char *hostname;
```

**DESCRIPTION**

These routines are useful for mapping 48 bit Ethernet numbers to their ASCII representations or their corresponding host names, and vice versa.

The function **ether\_ntoa()** converts a 48 bit Ethernet number pointed to by *e* to its standard ASCII representation; it returns a pointer to the ASCII string. The representation is of the form: *x:x:x:x:x:x* where *x* is a hexadecimal number between 0 and ff. The function **ether\_aton()** converts an ASCII string in the standard representation back to a 48 bit Ethernet number; the function returns NULL if the string cannot be scanned successfully.

The function **ether\_ntohost()** maps an Ethernet number (pointed to by *e*) to its associated hostname. The string pointed to by **hostname** must be long enough to hold the hostname and a null character. The function returns zero upon success and non-zero upon failure. Inversely, the function **ether\_hostton()** maps a hostname string to its corresponding Ethernet number; the function modifies the Ethernet number pointed to by *e*. The function also returns zero upon success and non-zero upon failure.

The function **ether\_line()** scans a line (pointed to by *l*) and sets the hostname and the Ethernet number (pointed to by *e*). The string pointed to by **hostname** must be long enough to hold the hostname and a null character. The function returns zero upon success and non-zero upon failure. The format of the scanned line is described by **ethers(5)**.

**FILES**

**/etc/ethers** (or the Network Information Service (NIS) maps **ethers.byaddr** and **ethers.byname**)

**SEE ALSO**

**ethers(5)**

**NOTES**

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed.

## NAME

execl, execl, execl, execlp, execlp – execute a file

## SYNOPSIS

```
int execl(path, arg0 [ , arg1, ... , argn ] (char *)0)
char *path, *arg0, *arg1, ..., *argn;

int execlv(path, argv)
char *path, *argv[ ];

int execl(path, arg0 [ , arg1, ... , argn ] (char *)0, envp)
char *path, *arg0, *arg1, ..., *argn, *envp[ ];

int execlp(file, arg0 [ , arg1, ... , argn ] (char *)0)
char *file, *arg0, *arg1, ..., *argn;

int execlv(file, argv)
char *file, *argv[ ];

extern char **environ;
```

## DESCRIPTION

These routines provide various interfaces to the `execve()` system call. Refer to `execve(2V)` for a description of their properties; only brief descriptions are provided here.

`exec()` in all its forms overlays the calling process with the named file, then transfers to the entry point of the core image of the file. There can be no return from a successful `exec()`; the calling core image is lost.

The *filename* argument is a pointer to the name of the file to be executed. The pointers `arg[0]`, `arg[1]`... address null-terminated strings. Conventionally `arg[0]` is the name of the file.

Two interfaces are available. `execl()` is useful when a known file with known arguments is being called; the arguments to `execl()` are the character strings constituting the file and the arguments; the first argument is conventionally the same as the file name (or its last component). A `(char *)0` argument must end the argument list. The cast to type `char *` insures portability.

The `execlv()` version is useful when the number of arguments is unknown in advance; the arguments to `execlv()` are the name of the file to be executed and a vector of strings containing the arguments. The last argument string must be followed by a 0 pointer.

When a C program is executed, it is called as follows:

```
main(argc, argv, envp)
int argc;
char **argv, **envp;
```

where *argc* is the argument count and *argv* is an array of character pointers to the arguments themselves. As indicated, *argc* is conventionally at least one and the first member of the array points to a string containing the name of the file.

*argv* is directly usable in another `execlv()` because `argv[argc]` is 0.

*envp* is a pointer to an array of strings that constitute the *environment* of the process. Each string consists of a name, an '=', and a null-terminated value. The array of pointers is terminated by a NULL pointer. The shell `sh(1)` passes an environment entry for each global shell variable defined when the program is called. See `environ(5V)` for some conventionally used names. The C run-time start-off routine places a copy of *envp* in the global cell *environ*, which is used by `execlv()` and `execl()` to pass the environment to any sub-programs executed by the current program.

`execlp()` and `execlvp()` are called with the same arguments as `execl()` and `execlv()`, but duplicate the shell's actions in searching for an executable *file* in a list of directories. The directory list is obtained from the environment.



**RETURN VALUES**

These functions return to the calling process only on failure. They return `-1` and set `errno` to indicate the error if *path* or *file* cannot be found, if it is not executable, if it does not start with a valid magic number (see `a.out(5)`), if maximum memory is exceeded, or if the arguments require too much space. Even for the super-user, at least one of the execute-permission bits must be set for a file to be executed.

**ERRORS**

If any of the following conditions occur, these functions will return and set `errno` to one of the following:

- |                     |   |
|---------------------|---|
| <b>E2BIG</b>        | The number of bytes used by the new process image's argument list and environment list is greater than <code>{ARG_MAX}</code> bytes (see <code>sysconf(2V)</code> ).  |
| <b>EACCES</b>       | Search permission is denied for a directory listed in the new process image file's path prefix.<br>The new process image file denies execution permission.<br>The new process image file is not a regular file.   |
| <b>ENAMETOOLONG</b> | The length of the <i>path</i> or <i>file</i> , or an element of the environment variable <code>PATH</code> prefixed to a file, exceeds <code>{PATH_MAX}</code> .<br>A pathname component is longer than <code>{NAME_MAX}</code> while <code>{_POSIX_NO_TRUNC}</code> is in effect for that file (see <code>pathconf(2V)</code> ). |
| <b>ENOENT</b>       | One or more components of the new process image file's pathname do not exist.   |
| <b>ENOTDIR</b>      | A component of the new process image file's path prefix is not a directory.   |

if the following condition occurs, `execl()`, `execv()`, and `execle()` set `errno` to:

- |                |  |
|----------------|--|
| <b>ENOEXEC</b> | The new process image file has the appropriate access permission, but is not in the proper format. |
|----------------|--|

If the following condition is detected, the exec functions set `errno` to:

- |               |  |
|---------------|--|
| <b>ENOMEM</b> | The new process image requires more memory than there is swap space available.<br>On Sun-3 systems, the new process image requires more than $2^{31}$ bytes. |
|---------------|--|

**SYSTEM V ERRORS**

In addition to the above, if the following condition occurs, the exec functions set `errno` to:

- |               |   |
|---------------|---|
| <b>ENOENT</b> | <i>path</i> or <i>file</i> points to a null pathname. |
|---------------|---|

**FILES**

- |                          |   |
|--------------------------|---|
| <code>/usr/bin/sh</code> | shell, invoked if command <i>file</i> found by <code>execlp()</code> or <code>execvp()</code> |
|--------------------------|---|

**SEE ALSO**

`csh(1)`, `sh(1)`, `execve(2V)`, `fork(2V)`, `pathconf(2V)`, `sysconf(2V)`, `a.out(5)`, `environ(5V)`

*Programming Utilities and Libraries*

**NAME**

**exit** – terminate a process after performing cleanup

**SYNOPSIS**

```
void  
exit(status)  
int status;
```

**DESCRIPTION**

**exit()** terminates a process by calling **exit(2V)** after calling any termination handlers named by calls to **on\_exit**. Normally, this is just the Standard I/O library function **\_cleanup**. **exit()** never returns.

**SEE ALSO**

**exit(2V)**, **intro(3)**, **on\_exit(3)**

## NAME

exportent, getexportent, setexportent, addexportent, remexportent, endexportent, getexportopt – get exported file system information

## SYNOPSIS

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

FILE *setexportent()

struct exportent *getexportent(filep)
FILE *filep;

int addexportent(filep, dirname, options)
FILE *filep;
char *dirname;
char *options;

int remexportent(filep, dirname)
FILE *filep;
char *dirname;

char *getexportopt(xent, opt)
struct exportent *xent;
char *opt;

void endexportent(filep)
FILE *filep;
```

## DESCRIPTION

These routines access the exported filesystem information in */etc/xtab*.

*setexportent()* opens the export information file and returns a file pointer to use with *getexportent*, *addexportent*, *remexportent*, and *endexportent*. *getexportent()* reads the next line from *filep* and returns a pointer to an object with the following structure containing the broken-out fields of a line in the file, */etc/xtab*. The fields have meanings described in *exports(5)*.

```
#define ACCESS_OPT "access" /* machines that can mount fs */
#define ROOT_OPT "root" /* machines with root access of fs */
#define RO_OPT "ro" /* export read-only */
#define ANON_OPT "anon" /* uid for anonymous requests */
#define SECURE_OPT "secure" /* require secure NFS for access */
#define WINDOW_OPT "window" /* expiration window for credential */

struct exportent {
    char *xent_dirname; /* directory (or file) to export */
    char *xent_options; /* options, as above */
};
```

*addexportent()* adds the *exportent()* to the end of the open file *filep*. It returns 0 if successful and -1 on failure. *remexportent()* removes the indicated entry from the list. It also returns 0 on success and -1 on failure. *getexportopt()* scans the *xent\_options* field of the *exportent()* structure for a substring that matches *opt*. It returns the string value of *opt*, or NULL if the option is not found.

*endexportent()* closes the file.

## FILES

```
/etc/exports
/etc/xtab
```

**SEE ALSO**

**exports(5), exportfs(8)**

**DIAGNOSTICS**

NULL pointer (0) returned on EOF or error.

**BUGS**

The returned **exportent()** structure points to static information that is overwritten in each call.

**NAME**

**fclose, fflush** – close or flush a stream

**SYNOPSIS**

**#include <stdio.h>**

**fclose(stream)**

**FILE \*stream;**

**fflush(stream)**

**FILE \*stream;**

**DESCRIPTION**

**fclose()** writes out any buffered data for the named stream, and closes the named stream. Buffers allocated by the standard input/output system are freed.

**fclose()** is performed automatically for all open files upon calling **exit(3)**.

**fflush()** writes any unwritten data for an output stream or an update stream in which the most recent operation was not input to be delivered to the host environment to the file; otherwise it is ignored. The named stream remains open.

**SYSTEM V DESCRIPTION**

When **fflush()** is called on a stream opened for reading, any unread data buffered in the stream is invalidated. When **fflush()** is called on a stream opened for reading, if the file is not already at EOF, and the file is one capable of seeking, the file offset of the underlying open file description is adjusted so the next operation on the open file description deals with the byte after the last byte read from or written to the stream being flushed.

**RETURN VALUES**

**fclose()** and **fflush()** return:

0 on success.

EOF if any error (such as trying to write to a file that has not been opened for writing) was detected.

**SEE ALSO**

**close(2V), exit(3), fopen(3V), setbuf(3V)**

**NAME**

error, feof, clearerr, fileno – stream status inquiries

**SYNOPSIS**

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

```
error(stream)
```

```
FILE *stream;
```

```
feof(stream)
```

```
FILE *stream;
```

```
clearerr(stream)
```

```
FILE *stream;
```

```
fileno(stream)
```

```
FILE *stream;
```

**DESCRIPTION**

**error()** returns non-zero when an error has occurred reading from or writing to the named stream, otherwise zero. Unless cleared by **clearerr()**, the error indication lasts until the stream is closed.

**feof()** returns non-zero when EOF has previously been detected reading the named input stream, otherwise zero. Unless cleared by **clearerr()**, the EOF indication lasts until the stream is closed.

**clearerr()** resets the error indication and EOF indication to zero on the named stream.

**fileno()** returns the integer file descriptor associated with the stream (see **open(2V)**).

**SYSTEM V DESCRIPTION**

**feof()** returns non-zero when EOF has previously been detected reading the named input stream, otherwise zero. Unless cleared by **clearerr()**, the EOF indication lasts until the stream is closed, however, operations which attempt to read from the stream will ignore the current state of the EOF indication and attempt to read from the file descriptor associated with the stream.

**SEE ALSO**

**open(2V)**, **fopen(3V)**

**NOTES**

These functions are defined in the C library and are also defined as macros in **<stdio.h>**.

## NAME

single\_to\_decimal, double\_to\_decimal, extended\_to\_decimal – convert floating-point value to decimal record

## SYNOPSIS

```
#include <floatingpoint.h>

void single_to_decimal(px, pm, pd, ps)
single *px ;
decimal_mode *pm;
decimal_record *pd;
fp_exception_field_type *ps;

void double_to_decimal(px, pm, pd, ps)
double *px ;
decimal_mode *pm;
decimal_record *pd;
fp_exception_field_type *ps;

void extended_to_decimal(px, pm, pd, ps)
extended *px ;
decimal_mode *pm;
decimal_record *pd;
fp_exception_field_type *ps;
```

## DESCRIPTION

The `floating_to_decimal()` functions convert the floating-point value at `*px` into a decimal record at `*pd`, observing the modes specified in `*pm` and setting exceptions in `*ps`. If there are no IEEE exceptions, `*ps` will be zero.

If `*px` is zero, infinity, or NaN, then only `pd->sign` and `pd->fpclass` are set. Otherwise `pd->exponent` and `pd->ds` are also set so that

$$(pd->sign)*(pd->ds)*10^{(pd->exponent)}$$

is a correctly rounded approximation to `*px`. `pd->ds` has at least one and no more than `DECIMAL_STRING_LENGTH-1` significant digits because one character is used to terminate the string with a null character.

`pd->ds` is correctly rounded according to the IEEE rounding modes in `pm->rd`. `*ps` has `fp_inexact` set if the result was inexact, and has `fp_overflow` set if the string result does not fit in `pd->ds` because of the limitation `DECIMAL_STRING_LENGTH`.

If `pm->df == floating_form`, then `pd->ds` always contains `pm->ndigits` significant digits. Thus if `*px == 12.34` and `pm->ndigits == 8`, then `pd->ds` will contain 12340000 and `pd->exponent` will contain -6.

If `pm->df == fixed_form` and `pm->ndigits >= 0`, then `pd->ds` always contains `pm->ndigits` after the point and as many digits as necessary before the point. Since the latter is not known in advance, the total number of digits required is returned in `pd->ndigits`; if that number `>= DECIMAL_STRING_LENGTH`, then `ds` is undefined. `pd->exponent` always gets `-pm->ndigits`. Thus if `*px == 12.34` and `pm->ndigits == 1`, then `pd->ds` gets 123, `pd->exponent` gets -1, and `pd->ndigits` gets 3.

If `pm->df == fixed_form` and `pm->ndigits < 0`, then `pd->ds` always contains `-pm->ndigits` trailing zeros; in other words, rounding occurs `-pm->ndigits` to the left of the decimal point, but the digits rounded away are retained as zeros. The total number of digits required is in `pd->ndigits`. `pd->exponent` always gets 0. Thus if `*px == 12.34` and `pm->ndigits == -1`, then `pd->ds` gets 10, `pd->exponent` gets 0, and `pd->ndigits` gets 2.

*pd->more* is not used.

**econvert()**, **fconvert()** and **gconvert()** (see **econvert(3)**), and **printf()** and **sprintf()** (see **printf(3V)**) all use **double\_to\_decimal()**.

**SEE ALSO**

**econvert(3)**, **printf(3V)**



**NAME**

floatingpoint – IEEE floating point definitions

**SYNOPSIS**

```
#include <sys/ieeefp.h>
#include <floatingpoint.h>
```

**DESCRIPTION**

This file defines constants, types, variables, and functions used to implement standard floating point according to ANSI/IEEE Std 754-1985. The variables and functions are implemented in `libc.a`. The included file `<sys/ieeefp.h>` defines certain types of interest to the kernel.

**IEEE Rounding Modes:**

- fp\_direction\_type** The type of the IEEE rounding direction mode. Note: the order of enumeration varies according to hardware.
- fp\_direction** The IEEE rounding direction mode currently in force. This is a global variable that is intended to reflect the hardware state, so it should only be written indirectly through a function like `ieee_flags("set","direction",...)` that also sets the hardware state.
- fp\_precision\_type** The type of the IEEE rounding precision mode, which only applies on systems that support extended precision such as Sun-3 systems with 68881's.
- fp\_precision** The IEEE rounding precision mode currently in force. This is a global variable that is intended to reflect the hardware state on systems with extended precision, so it should only be written indirectly through a function like `ieee_flags("set","precision",...)`.

**SIGFPE handling:**

- sigfpe\_code\_type** The type of a SIGFPE code.
- sigfpe\_handler\_type** The type of a user-definable SIGFPE exception handler called to handle a particular SIGFPE code.
- SIGFPE\_DEFAULT** A macro indicating the default SIGFPE exception handling, namely to perform the exception handling specified by calls to `ieee_handler(3M)`, if any, and otherwise to dump core using `abort(3)`.
- SIGFPE\_IGNORE** A macro indicating an alternate SIGFPE exception handling, namely to ignore and continue execution.
- SIGFPE\_ABORT** A macro indicating an alternate SIGFPE exception handling, namely to abort with a core dump.

**IEEE Exception Handling:**

- N\_IEEE\_EXCEPTION** The number of distinct IEEE floating-point exceptions.
- fp\_exception\_type** The type of the `N_IEEE_EXCEPTION` exceptions. Each exception is given a bit number.
- fp\_exception\_field\_type** The type intended to hold at least `N_IEEE_EXCEPTION` bits corresponding to the IEEE exceptions numbered by `fp_exception_type`. Thus `fp_inexact` corresponds to the least significant bit and `fp_invalid` to the fifth least significant bit. Note: some operations may set more than one exception.
- fp\_accrued\_exceptions** The IEEE exceptions between the time this global variable was last cleared, and the last time a function like `ieee_flags("get","exception",...)` was called to update the variable by obtaining the hardware state.

**ieee\_handlers** An array of user-specifiable signal handlers for use by the standard SIGFPE handler for IEEE arithmetic-related SIGFPE codes. Since IEEE trapping modes correspond to hardware modes, elements of this array should only be modified with a function like **ieee\_handler(3M)** that performs the appropriate hardware mode update. If no **sigfpe\_handler** has been declared for a particular IEEE-related SIGFPE code, then the related **ieee\_handlers** will be invoked.

**IEEE Formats and Classification:**

**single;extended** Definitions of IEEE formats.

**fp\_class\_type** An enumeration of the various classes of IEEE values and symbols.

**IEEE Base Conversion:**

The functions described under **floating\_to\_decimal(3)** and **decimal\_to\_floating(3)** not only satisfy the IEEE Standard, but also the stricter requirements of correct rounding for all arguments.

**DECIMAL\_STRING\_LENGTH**

The length of a **decimal\_string**.

**decimal\_string** The digit buffer in a **decimal\_record**.

**decimal\_record** The canonical form for representing an unpacked decimal floating-point number.

**decimal\_form** The type used to specify fixed or floating binary to decimal conversion.

**decimal\_mode** A struct that contains specifications for conversion between binary and decimal.

**decimal\_string\_form** An enumeration of possible valid character strings representing floating-point numbers, infinities, or NaNs.

**SEE ALSO**

**abort(3)**, **decimal\_to\_floating(3)**, **econvert(3)**, **floating\_to\_decimal(3)**, **ieee\_flags(3M)**, **ieee\_handler(3M)**, **sigfpe(3)**, **string\_to\_decimal(3)**, **strtod(3)**

**NAME**

`fopen`, `freopen`, `fdopen` – open a stream

**SYNOPSIS**

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

```
FILE *fopen(filename, type)
```

```
char *filename, *type;
```

```
FILE *freopen(filename, type, stream)
```

```
char *filename, *type;
```

```
FILE *stream;
```

```
FILE *fdopen(fd, type)
```

```
int fd;
```

```
char *type;
```

**DESCRIPTION**

`fopen()` opens the file named by *filename* and associates a stream with it. If the open succeeds, `fopen()` returns a pointer to be used to identify the stream in subsequent operations.

*filename* points to a character string that contains the name of the file to be opened.

*type* is a character string having one of the following values:

<code>r</code>	open for reading
<code>w</code>	truncate or create for writing
<code>a</code>	append: open for writing at end of file, or create for writing
<code>r+</code>	open for update (reading and writing)
<code>w+</code>	truncate or create for update
<code>a+</code>	append; open or create for update at EOF

`freopen()` opens the file named by *filename* and associates the stream pointed to by *stream* with it. The *type* argument is used just as in `fopen`. The original stream is closed, regardless of whether the open ultimately succeeds. If the open succeeds, `freopen()` returns the original value of *stream*.

`freopen()` is typically used to attach the preopened streams associated with `stdin`, `stdout`, and `stderr` to other files.

`fdopen()` associates a stream with the file descriptor *fd*. File descriptors are obtained from calls like `open(2V)`, `dup(2V)`, `creat(2V)`, or `pipe(2V)`, which open files but do not return streams. Streams are necessary input for many of the Section 3S library routines. The *type* of the stream must agree with the access permissions of the open file.

When a file is opened for update, both input and output may be done on the resulting stream. However, output may not be directly followed by input without an intervening `fseek(3S)` or `rewind()`, and input may not be directly followed by output without an intervening `fseek()`, `rewind()`, or an input operation which encounters EOF.

When a file is opened for update, both input and output may be done on the resulting stream. However, output may not be directly followed by input without an intervening `fseek()` or `rewind()`, and input may not be directly followed by output without an intervening `fseek()`, `rewind()`, or an input operation which encounters end-of-file.

**SYSTEM V DESCRIPTION**

When a file is opened for append (that is, when *type* is *a* or *a+*), it is impossible to overwrite information already in the file. `fseek()` may be used to reposition the file pointer to any position in the file, but when output is written to the file, the current file pointer is disregarded. All output is written at the end of the file and causes the file pointer to be repositioned at the end of the output. If two separate processes open the same file for append, each process may write freely to the file without fear of destroying output being written by the other. The output from the two processes will be intermixed in the file in the order in which it is written.

**RETURN VALUES**

On success, `fopen()`, `freopen()`, and `fdopen()` return a pointer to `FILE` which identifies the opened stream. On failure, they return `NULL`.

**SEE ALSO**

`open(2V)`, `pipe(2V)`, `fclose(3V)`, `fseek(3S)`

**BUGS**

In order to support the same number of open files that the system does, `fopen()` must allocate additional memory for data structures using `calloc()` after 64 files have been opened. This confuses some programs which use their own memory allocators.

**NAME**

**fread, fwrite** – buffered binary input/output

**SYNOPSIS**

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

```
int fread (ptr, size, nitems, stream)
```

```
char *ptr;
```

```
int size;
```

```
int nitems;
```

```
FILE *stream;
```

```
int fwrite (ptr, size, nitems, stream)
```

```
char *ptr;
```

```
int size;
```

```
int nitems;
```

```
FILE *stream;
```

**DESCRIPTION**

**fread()** reads, into a block pointed to by *ptr*, *nitems* items of data from the named input stream *stream*, where an item of data is a sequence of bytes (not necessarily terminated by a null byte) of length *size*. It returns the number of items actually read. **fread()** stops reading if an end-of-file or error condition is encountered while reading from *stream*, or if *nitems* items have been read. **fread()** leaves the file pointer in *stream*, if defined, pointing to the byte following the last byte read if there is one. **fread()** does not change the contents of the file referred to by *stream*.

**fwrite()** writes at most *nitems* items of data from the block pointed to by *ptr* to the named output stream *stream*. It returns the number of items actually written. **fwrite()** stops writing when it has written *nitems* items of data or if an error condition is encountered on *stream*. **fwrite()** does not change the contents of the block pointed to by *ptr*.

If *size* or *nitems* is non-positive, no characters are read or written and 0 is returned by both **fread()** and **fwrite()**.

**SEE ALSO**

**read(2V), write(2V), fopen(3V), getc(3V), gets(3S), putc(3S), puts(3S), printf(3V), scanf(3V)**

**DIAGNOSTICS**

**fread()** and **fwrite()** return 0 upon end of file or error.

**NAME**

**fseek, ftell, rewind** – reposition a stream

**SYNOPSIS**

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

```
fseek(stream, offset, ptrname)
```

```
FILE *stream;
```

```
long offset;
```

```
long ftell(stream)
```

```
FILE *stream;
```

```
rewind(stream)
```

```
FILE *stream;
```

**DESCRIPTION**

**fseek()** sets the position of the next input or output operation on the stream. The new position is at the signed distance *offset* bytes from the beginning, the current position, or the end of the file, according as *ptrname* has the value 0, 1, or 2.

**rewind(stream)** is equivalent to **fseek(stream, 0L, 0)**, except that no value is returned.

**fseek()** and **rewind()** undo any effects of **ungetc(3S)**.

After **fseek()** or **rewind()**, the next operation on a file opened for update may be either input or output.

**ftell()** returns the offset of the current byte relative to the beginning of the file associated with the named stream.

**SEE ALSO**

**lseek(2V), fopen(3V), popen(3S), ungetc(3S)**

**DIAGNOSTICS**

**fseek()** returns **-1** for improper seeks, otherwise zero. An improper seek can be, for example, an **fseek()** done on a file associated with a non-seekable device, such as a tty or a pipe; in particular, **fseek()** may not be used on a terminal, or on a file opened using **popen(3S)**.

**WARNING**

Although on the UNIX system an offset returned by **ftell()** is measured in bytes, and it is permissible to seek to positions relative to that offset, portability to a (non-UNIX) system requires that an offset be used by **fseek()** directly. Arithmetic may not meaningfully be performed on such an offset, which is not necessarily measured in bytes.

**NAME**

**ftok** – standard interprocess communication package

**SYNOPSIS**

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

```
#include <sys/ipc.h>
```

```
key_t ftok(path, id)
```

```
char *path;
```

```
char id;
```

**DESCRIPTION**

All interprocess communication facilities require the user to supply a key to be used by the **msgget(2)**, **semget(2)**, and **shmget(2)** system calls to obtain interprocess communication identifiers. One suggested method for forming a key is to use the **ftok()** subroutine described below. Another way to compose keys is to include the project ID in the most significant byte and to use the remaining portion as a sequence number. There are many other ways to form keys, but it is necessary for each system to define standards for forming them. If some standard is not adhered to, it will be possible for unrelated processes to unintentionally interfere with each other's operation. Therefore, it is strongly suggested that the most significant byte of a key in some sense refer to a project so that keys do not conflict across a given system.

**ftok()** returns a key based on *path* and ID that is usable in subsequent **msgget**, **semget**, and **shmget()** system calls. *path* must be the path name of an existing file that is accessible to the process. ID is a character which uniquely identifies a project. Note: **ftok()** will return the same key for linked files when called with the same ID and that it will return different keys when called with the same file name but different IDs.

**SEE ALSO**

**intro(2)**, **msgget(2)**, **semget(2)**, **shmget(2)**

**DIAGNOSTICS**

**ftok()** returns (**key\_t**) **-1** if *path* does not exist or if it is not accessible to the process.

**WARNING**

If the file whose *path* is passed to **ftok()** is removed when keys still refer to the file, future calls to **ftok()** with the same *path* and ID will return an error. If the same file is recreated, then **ftok()** is likely to return a different key than it did the original time it was called.

**NAME**

`ftw` – walk a file tree

**SYNOPSIS**

```
#include <ftw.h>

int ftw(path, fn, depth)
char *path;
int (*fn)();
int depth;
```

**DESCRIPTION**

`ftw()` recursively descends the directory hierarchy rooted in *path*. For each object in the hierarchy, `ftw()` calls *fn*, passing it a pointer to a null-terminated character string containing the name of the object, a pointer to a `stat()` structure (see `stat(2V)`) containing information about the object, and an integer. Possible values of the integer, defined in the `<ftw.h>` header file, are `FTW_F` for a file, `FTW_D` for a directory, `FTW_DNR` for a directory that cannot be read, and `FTW_NS` for an object for which `stat()` could not successfully be executed. If the integer is `FTW_DNR`, descendants of that directory will not be processed. If the integer is `FTW_NS`, the `stat()` structure will contain garbage. An example of an object that would cause `FTW_NS` to be passed to *fn* would be a file in a directory with read but without execute (search) permission.

`ftw()` visits a directory before visiting any of its descendants.

The tree traversal continues until the tree is exhausted, an invocation of *fn* returns a nonzero value, or some error is detected within `ftw()` (such as an I/O error). If the tree is exhausted, `ftw()` returns zero. If *fn* returns a nonzero value, `ftw()` stops its tree traversal and returns whatever value was returned by *fn*. If `ftw()` detects an error, it returns `-1`, and sets the error type in `errno`.

`ftw()` uses one file descriptor for each level in the tree. The *depth* argument limits the number of file descriptors so used. If *depth* is zero or negative, the effect is the same as if it were 1. *depth* must not be greater than the number of file descriptors currently available for use. `ftw()` will run more quickly if *depth* is at least as large as the number of levels in the tree.

**SEE ALSO**

`stat(2V)`, `malloc(3V)`

**BUGS**

Because `ftw()` is recursive, it is possible for it to terminate with a memory fault when applied to very deep file structures.

It could be made to run faster and use less storage on deep structures at the cost of considerable complexity.

`ftw()` uses `malloc(3V)` to allocate dynamic storage during its operation. If `ftw()` is forcibly terminated, such as by `longjmp()` being executed by *fn* or an interrupt routine, `ftw()` will not have a chance to free that storage, so it will remain permanently allocated. A safe way to handle interrupts is to store the fact that an interrupt has occurred, and arrange to have *fn* return a nonzero value at its next invocation.



**NAME**

getacinfo, getacdir, getacflg, getacmin, setac, endac – get audit control file information

**SYNOPSIS**

```
int getacdir(dir, len)
char *dir;
int len;

int getacmin(min_val)
int *min_val;

int getacflg(auditstring, len)
char *auditstring;
int len;

void setac()

void endac()
```

**DESCRIPTION**

When first called, **getacdir()** provides information about the first audit directory in the **audit\_control** file; thereafter, it returns the next directory in the file. Successive calls list all the directories listed in **audit\_control(5)**. The parameter *len* specifies the length of the buffer *dir*. On return, *dir* points to the directory entry.

**getacmin()** reads the minimum value from the **audit\_control** file and returns the value in *min\_val*. The minimum value specifies how full the file system to which the audit files are being written can get before the script **audit\_warn** is invoked.

**getacflg()** reads the system audit value from the **audit\_control** file and returns the value in *auditstring*. The parameter *len* specifies the length of the buffer *auditstring*.

Calling **setac** rewinds the **audit\_control** file to allow repeated searches.

Calling **endac** closes the **audit\_control** file when processing is complete.

**RETURN VALUES**

**getacdir()**, **getacflg()** and **getacmin()** return:

- 0 on success.
- 2 on failure and set **errno** to indicate the error.

**getacmin()** and **getacflg()** return:

- 1 on EOF.

**getacdir()** returns:

- 1 on EOF.
- 2 if the directory search had to start from the beginning because one of the other functions was called between calls to **getacdir()**.

These functions return:

- 3 if the directory entry format in the **audit\_control** file is incorrect.

**getacdir()** and **getacflg()** return:

- 3 if the input buffer is too short to accommodate the record.

**SEE ALSO**

**audit\_control(5)**

**NAME**

getauditflagsbin, getauditflagschar – convert audit flag specifications

**SYNOPSIS**

```
#include <sys/label.h>
#include <sys/audit.h>
#include <sys/aevents.h>

int getauditflagsbin(auditstring, masks)
char *auditstring;
audit_state_t *masks;

int getauditflagschar(auditstring, masks, verbose)
char *auditstring;
audit_state_t *masks;
int verbose;
```

**DESCRIPTION**

**getauditflagsbin()** converts the character representation of audit values pointed to by *auditstring* into **audit\_state\_t** fields pointed to by *masks*. These fields indicate which events are to be audited when they succeed and which are to be audited when they fail. The character string syntax is described in **audit\_control(5)**.

**getauditflagschar()** converts the **audit\_state\_t** fields pointed to by *masks* into a string pointed to by *auditstring*. If *verbose* is zero, the short (2-character) flag names are used. If *verbose* is non-zero, the long flag names are used. *auditstring* should be large enough to contain the ASCII representation of the events.

*auditstring* contains a series of event names, each one identifying a single audit class, separated by commas. The **audit\_state\_t** fields pointed to by *masks* correspond to binary values defined in *audit.h*.

**DIAGNOSTICS**

-1 is returned on error and 0 on success.

**SEE ALSO**

**audit.log(5)**, **audit\_control(5)**

**BUGS**

This is not a very extensible interface.

**NAME**

`getc`, `getchar`, `fgetc`, `getw` – get character or integer from stream

**SYNOPSIS**

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

```
int getc(stream)
```

```
FILE *stream;
```

```
int getchar()
```

```
int fgetc(stream)
```

```
FILE *stream;
```

```
int getw(stream)
```

```
FILE *stream;
```

**DESCRIPTION**

`getc()` returns the next character (that is, byte) from the named input stream, as an integer. It also moves the file pointer, if defined, ahead one character in stream. `getchar()` is defined as `getc(stdin)`. `getc()` and `getchar()` are macros.

`fgetc()` behaves like `getc()`, but is a function rather than a macro. `fgetc()` runs more slowly than `getc()`, but it takes less space per invocation and its name can be passed as an argument to a function.

`getw()` returns the next C `int` (*word*) from the named input stream. `getw()` increments the associated file pointer, if defined, to point to the next word. The size of a word is the size of an integer and varies from machine to machine. `getw()` assumes no special alignment in the file.

**RETURN VALUES**

On success, `getc()`, `getchar()` and `fgetc()` return the next character from the named input stream as an integer. On failure, or on EOF, they return EOF. The EOF condition is remembered, even on a terminal, and all subsequent operations which attempt to read from the stream will return EOF until the condition is cleared with `clearerr()` (see `ferror(3V)`).

`getw()` returns the next C `int` from the named input stream on success. On failure, or on EOF, it returns EOF, but since EOF is a valid integer, use `ferror(3V)` to detect `getw()` errors.

**SYSTEM V RETURN VALUES**

On failure, or on EOF, these functions return EOF. The EOF condition is remembered, even on a terminal, however, operations which attempt to read from the stream will ignore the current state of the EOF indication and attempt to read from the file descriptor associated with the stream.

**SEE ALSO**

`ferror(3V)`, `fopen(3V)`, `fread(3S)`, `gets(3S)`, `putc(3S)`, `scanf(3V)`, `ungetc(3S)`

**WARNINGS**

If the integer value returned by `getc()`, `getchar()`, or `fgetc()` is stored into a character variable and then compared against the integer constant EOF, the comparison may never succeed, because sign-extension of a character on widening to integer is machine-dependent.

**BUGS**

Because it is implemented as a macro, `getc()` treats a stream argument with side effects incorrectly. In particular, `getc(*f++)` does not work sensibly. `fgetc()` should be used instead.

Because of possible differences in word length and byte ordering, files written using `putw()` are machine-dependent, and may not be readable using `getw()` on a different processor.

**NAME**

`getcwd` – get pathname of current working directory

**SYNOPSIS**

```
char *getcwd(buf, size)
char *buf;
int size;
```

**DESCRIPTION**

`getcwd()` returns a pointer to the current directory pathname. The value of *size* must be at least two greater than the length of the pathname to be returned.

If *buf* is a NULL pointer, `getcwd()` will obtain *size* bytes of space using `malloc(3V)`. In this case, the pointer returned by `getcwd()` may be used as the argument in a subsequent call to `free()`.

The function is implemented by using `popen(3S)` to pipe the output of the `pwd(1)` command into the specified string space.

**RETURN VALUES**

`getcwd()` returns a pointer to the current directory pathname on success. If *size* is not large enough, or if an error occurs in a lower-level function, `getcwd()` returns NULL and sets `errno` to indicate the error.

**ERRORS**

`EINVAL`            *size* is less than or equal to zero.  
`ERANGE`            *size* is greater than zero, but is smaller than the length of the pathname plus 1.

If the following condition is detected, `getcwd()` sets `errno` to:

`EACCES`            Read or search permission is denied for a component of the pathname.

**EXAMPLES**

```
char *cwd, *getcwd();
.
.
.
if ((cwd = getcwd((char *)NULL, 64)) == NULL) {
    perror("pwd");
    exit(1);
}
printf("%s\n", cwd);
```

**SEE ALSO**

`pwd(1)`, `getwd(3)`, `malloc(3V)`, `popen(3S)`

**BUGS**

Since this function uses `popen()` to create a pipe to the `pwd` command, it is slower than `getwd()` and gives poorer error diagnostics. `getcwd()` is provided only for compatibility with other UNIX operating systems.

**NAME**

getenv – return value for environment name

**SYNOPSIS**

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

```
char *getenv(name)
```

```
char *name;
```

**DESCRIPTION**

**getenv()** searches the environment list (see **environ(5V)**) for a string of the form *name=value*, and returns a pointer to the string *value* if such a string is present. Otherwise, **getenv()** returns NULL.

**RETURN VALUES**

On success, **getenv()** returns a pointer to a string containing the value for the specified *name*. If the specified *name* cannot be found, it returns NULL.

**SEE ALSO**

**environ(5V)**, **execve(2V)**, **putenv(3)**

**NAME**

getfauditflags – generates the process audit state

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/audit.h>
#include <sys/label.h>

void getfauditflags(usremasks, usrdmasks, lastmasks)
audit_state_t *usremasks;
audit_state_t *usrdmasks;
audit_state_t *lastmasks;
```

**DESCRIPTION**

**getfauditflags** generates the process audit state from the user audit value as input to **getfauditflags** and the system audit value as specified in the **audit\_control** file. **getfauditflags** obtains the system audit value by calling **getacflg**. The user audit value, pointed to by *usremasks* and *usrdmasks* is passed into **getfauditflags**.

*usremasks* points to **audit\_state\_t** fields which contains two values. The first value defines which events are *always* to be audited when they succeed. The second value defines which events are always to be audited when they fail.

*usrdmasks* also points to **audit\_state\_t** fields which contains two values. The first value defines which events are *never* to be audited when they succeed. The second value defines which events are never to be audited when they fail.

The structures pointed to by *usremasks* and *usrdmasks* may be obtained from the **passwd.adjunct** file by calling **getpwaent()** which returns a pointer to a structure containing all **passwd.adjunct** fields for a user.

*lastmasks* points to **audit\_state\_t** as well. The first value defines which events are to be audited when they succeed and the second value defines which events are to be audited when they fail.

Both *usremasks* and *usrdmasks* override the values in the system audit values.

**DIAGNOSTICS**

-1 is returned on error and 0 on success.

**SEE ALSO**

**getauditflags(3)**, **getacinfo(3)**, **audit.log(5)**, **audit\_control(5)**

**NAME**

*getfsent*, *getfsspec*, *getfsfile*, *getfstype*, *setfsent*, *endfsent* – get file system descriptor file entry

**SYNOPSIS**

```
#include <fstab.h>

struct fstab *getfsent()
struct fstab *getfsspec(spec)
char *spec;

struct fstab *getfsfile(file)
char *file;

struct fstab *getfstype(type)
char *type;

int setfsent()

int endfsent()
```

**DESCRIPTION**

These routines are included for compatibility with 4.2 BSD; they have been superseded by the *getmntent(3)* library routines.

*getfsent*, *getfsspec*, *getfstype*, and *getfsfile* each return a pointer to an object with the following structure containing the broken-out fields of a line in the file system description file, *<fstab.h>*.

```
struct fstab {
    char *fs_spec;
    char *fs_file;
    char *fs_type;
    int fs_freq;
    int fs_passno;
};
```

The fields have meanings described in *fstab(5)*.

*getfsent()* reads the next line of the file, opening the file if necessary.

*getfsent()* opens and rewinds the file.

*endfsent* closes the file.

*getfsspec* and *getfsfile* sequentially search from the beginning of the file until a matching special file name or file system file name is found, or until EOF is encountered. *getfstype* does likewise, matching on the file system type field.

**FILES**

*/etc/fstab*

**SEE ALSO**

*fstab(5)*

**DIAGNOSTICS**

Null pointer (0) returned on EOF or error.

**BUGS**

The return value points to static information which is overwritten in each call.

**NAME**

getgraent, getgranam, setgraent, endgraent, fgetgraent -- get group adjunct file entry

**SYNOPSIS**

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

struct group_adjunct *getgraent()

struct group_adjunct *getgranam(name)
char *name;

struct group_adjunct *fgetgraent(f)
FILE *f;

void setgraent()

void endgraent()
```

**DESCRIPTION**

**getgraent()** and **getgranam()** each return pointers to an object with the following structure containing the broken-out fields of a line in the group adjunct file. Each line contains a **group\_adjunct** structure, defined in the **<grpadj.h>** header file.

```
struct group_adjunct {
    char *gra_name;    /* the name of the group */
    char *gra_passwd; /* the encrypted group password */
};
```

When first called, **getgraent()** returns a pointer to a **group\_adjunct** structure corresponding to the first line in the file. Thereafter, it returns a pointer to the next **group\_adjunct** structure in the file. So successive calls may be used to traverse the entire file.

For locating a particular group, **getgranam()** searches through the file until it finds group *filename*, then returns a pointer to that structure.

A call to **getgraent()** rewinds the group adjunct file to allow repeated searches. A call to **endgraent()** closes the group adjunct file when processing is complete.

Because read access is required on **/etc/security/group.adjunct**, **getgraent()** and **getgranam()** will fail unless the calling process has effective UID of root.

**FILES**

```
/etc/security/group.adjunct
/var/yp/domainname/group.adjunct
```

**SEE ALSO**

**getlogin(3V)**, **getgrent(3V)**, **getpwaent(3)**, **getpwent(3V)**, **ypserv(8)**

**DIAGNOSTICS**

A NULL pointer is returned on end-of-file or error.

**BUGS**

All information is contained in a static area, so it must be copied if it is to be saved.



**NAME**

getgrent, getgrgid, getgrnam, setgrent, endgrent, fgetgrent – get group file entry

**SYNOPSIS**

```
#include <grp.h>

struct group *getgrent()

struct group *getgrgid(gid)
int gid;

struct group *getgrnam(name)
char *name;

void setgrent()

void endgrent()

struct group *fgetgrent(f)
FILE *f;
```

**DESCRIPTION**

**getgrent()**, **getgrgid()** and **getgrnam()** each return pointers to an object with the following structure containing the fields of a line in the group file. Each line contains a “group” structure, defined in **<grp.h>**.

```
struct group {
    char   *gr_name;      /* name of the group */
    char   *gr_passwd;   /* encrypted password of the group */
    gid_t  gr_gid;       /* numerical group ID */
    char   **gr_mem;     /* null-terminated array of pointers to the
                           individual member names */
};
```

**getgrent()** when first called returns a pointer to the first group structure in the file; thereafter, it returns a pointer to the next group structure in the file; so, successive calls may be used to search the entire file. **getgrgid()** searches from the beginning of the file until a numerical group ID matching **gid** is found and returns a pointer to the particular structure in which it was found. **getgrnam()** searches from the beginning of the file until a group name matching **name** is found and returns a pointer to the particular structure in which it was found. If an end-of-file or an error is encountered on reading, these functions return a NULL pointer.

A call to **setgrent()** has the effect of rewinding the group file to allow repeated searches. **endgrent()** may be called to close the group file when processing is complete.

**fgetgrent()** returns a pointer to the next group structure in the stream **f**, which must refer to an open file in the same format as the group file **/etc/group**.

**RETURN VALUES**

**getgrent()**, **getgrgid()**, and **getgrnam()** return a pointer to **struct group** on success. On EOF or error, they return NULL.

**FILES**

**/etc/group**

**SEE ALSO**

**getlogin(3V)**, **getpwent(3V)**, **group(5)**, **ypserv(8)**

**BUGS**

All information is contained in a static area, so it must be copied if it is to be saved.

Unlike the corresponding routines for passwords (see **getpwent(3v)**), which always search the entire file, these routines start searching from the current file location.

**WARNING**

The above routines use the standard I/O library, which increases the size of programs not otherwise using standard I/O more than might be expected.

**NAME**

gethostent, gethostbyaddr, gethostbyname, sethostent, endhostent – get network host entry

**SYNOPSIS**

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

struct hostent *gethostent()

struct hostent *gethostbyname(name)
char *name;

struct hostent *gethostbyaddr(addr, len, type)
char *addr;
int len, type;

sethostent(stayopen)
int stayopen
endhostent()
```

**DESCRIPTION**

**gethostent**, **gethostbyname**, and **gethostbyaddr**() each return a pointer to an object with the following structure containing the broken-out fields of a line in the network host data base, */etc/hosts*. In the case of **gethostbyaddr**(), *addr* is a pointer to the binary format address of length *len* (not a character string).

```
struct hostent {
    char *h_name; /* official name of host */
    char **h_aliases; /* alias list */
    int h_addrtype; /* address type */
    int h_length; /* length of address */
    char **h_addr_list; /* list of addresses from name server */
};
```

The members of this structure are:

<b>h_name</b>	Official name of the host.
<b>h_aliases</b>	A zero terminated array of alternate names for the host.
<b>h_addrtype</b>	The type of address being returned; currently always AF_INET.
<b>h_length</b>	The length, in bytes, of the address.
<b>h_addr_list</b>	A pointer to a list of network addresses for the named host. Host addresses are returned in network byte order.

**gethostent**() reads the next line of the file, opening the file if necessary.

**sethostent**() opens and rewinds the file. If the *stayopen* flag is non-zero, the host data base will not be closed after each call to **gethostent**() (either directly, or indirectly through one of the other “gethost” calls).

**endhostent**() closes the file.

**gethostbyname**() and **gethostbyaddr**() sequentially search from the beginning of the file until a matching host name or host address is found, or until end-of-file is encountered. Host addresses are supplied in network order.

**FILES**

*/etc/hosts*

**SEE ALSO**

**hosts(5)**, **ypserv(8)**

**DIAGNOSTICS**

A NULL pointer is returned on end-of-file or error.

**BUGS**

All information is contained in a static area so it must be copied if it is to be saved. Only the Internet address format is currently understood.

**NAME**

getlogin – get login name

**SYNOPSIS**

**char \*getlogin()**

**DESCRIPTION**

**getlogin()** returns a pointer to the login name as found in **/etc/utmp**. It may be used in conjunction with **getpwnam()** to locate the correct password file entry when the same user ID is shared by several login names.

If **getlogin()** is called within a process that is not attached to a terminal, or if there is no entry in **/etc/utmp** for the process's terminal, it returns a NULL pointer. The correct procedure for determining the login name is to call **cuserid()**, or to call **getlogin()** and, if it fails, to call **getpwuid(getuid())**.

**FILES**

**/etc/utmp**

**SEE ALSO**

**cuserid(3v)**, **getpwent(3v)**, **utmp(5V)**

**RETURN VALUES**

**getlogin()** returns a pointer to the login name on success. If the name is not found, it returns NULL.

**BUGS**

The return values point to static data whose content is overwritten by each call.

**getlogin()** does not work for processes running under a **pty** (for example, emacs shell buffers, or shell tools) unless the program “fakes” the login name in the **/etc/utmp** file.

## NAME

getmntent, setmntent, addmntent, endmntent, hasmntopt – get file system descriptor file entry

## SYNOPSIS

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

FILE *setmntent(FILE *filep, char *type)
char *filep;
char *type;

struct mntent *getmntent(FILE *filep)
FILE *filep;

int addmntent(FILE *filep, struct mntent *mnt)
FILE *filep;
struct mntent *mnt;

char *hasmntopt(struct mntent *mnt, char *opt)
struct mntent *mnt;
char *opt;

int endmntent(FILE *filep)
FILE *filep;
```

## DESCRIPTION

These routines replace the `getfsent()` routines for accessing the file system description file `/etc/fstab`. They are also used to access the mounted file system description file `/etc/mntab`.

`setmntent()` opens a file system description file and returns a file pointer which can then be used with `getmntent`, `addmntent`, or `endmntent`. The `type` argument is the same as in `fopen(3V)`. `getmntent()` reads the next line from `filep` and returns a pointer to an object with the following structure containing the broken-out fields of a line in the file system description file, `<mntent.h>`. On failure, `getmntent()` returns the NULL pointer. The fields have meanings described in `fstab(5)`.

```
struct mntent{
    char *mnt_fstype; /* name of mounted file system */
    char *mnt_dir;    /* file system path prefix */
    char *mnt_type;   /* MNTTYPE_* */
    char *mnt_opts;   /* MNTOPT* */
    int mnt_freq;     /* dump frequency, in days */
    int mnt_passno;   /* pass number on parallel fsck */
};
```

`addmntent()` adds the `mntent` structure `mnt` to the end of the open file `filep`. `addmntent()` returns 0 on success, 1 on failure. Note: `filep` has to be opened for writing if this is to work. `hasmntopt()` scans the `mnt_opts` field of the `mntent` structure `mnt` for a substring that matches `opt`. It returns the address of the substring if a match is found, 0 otherwise. `endmntent()` closes the file. It always returns 1, so should be treated as type `void`.

## FILES

`/etc/fstab`  
`/etc/mntab`

## SEE ALSO

`fopen(3V)`, `getfsent(3)`, `fstab(5)`

## DIAGNOSTICS

NULL pointer (0) returned on EOF or error.

**BUGS**

The returned `mntent` structure points to static information that is overwritten in each call.

**NAME**

getnetent, getnetbyaddr, getnetbyname, setnetent, endnetent – get network entry

**SYNOPSIS**

```
#include <netdb.h>

struct netent *getnetent()

struct netent *getnetbyname(name)
char *name;

struct netent *getnetbyaddr(net, type)
long net;
int type;

setnetent (stayopen)
int stayopen;

endnetent()
```

**DESCRIPTION**

getnetent, getnetbyname, and getnetbyaddr() each return a pointer to an object with the following structure containing the broken-out fields of a line in the network data base, */etc/networks*.

```
struct netent {
    char    *n_name;        /* official name of net */
    char    **n_aliases;   /* alias list */
    int     n_addrtype;    /* net number type */
    long    n_net;         /* net number */
};
```

The members of this structure are:

<b>n_name</b>	The official name of the network.
<b>n_aliases</b>	A zero terminated list of alternate names for the network.
<b>n_addrtype</b>	The type of the network number returned; currently only AF_INET.
<b>n_net</b>	The network number. Network numbers are returned in machine byte order.

getnetent() reads the next line of the file, opening the file if necessary.

setnetent() opens and rewinds the file. If the *stayopen* flag is non-zero, the net data base will not be closed after each call to setnetent() (either directly, or indirectly through one of the other “getnet” calls).

endnetent() closes the file.

getnetbyname() and getnetbyaddr() sequentially search from the beginning of the file until a matching net name or net address and type is found, or until end-of-file is encountered. Network numbers are supplied in host order.

**FILES**

*/etc/networks*

**SEE ALSO**

networks(5), yperv(8)

**DIAGNOSTICS**

A NULL pointer is returned on end-of-file or error.

**BUGS**

All information is contained in a static area so it must be copied if it is to be saved.

Only Internet network numbers are currently understood.



**NAME**

getnetgrent, setnetgrent, endnetgrent, inetgr – get network group entry

**SYNOPSIS**

```
getnetgrent(machinep, userp, domainp)
char **machinep, **userp, **domainp;

setnetgrent(netgroup)
char *netgroup

endnetgrent()

inetgr(netgroup, machine, user, domain)
char *netgroup, *machine, *user, *domain;
```

**DESCRIPTION**

**getnetgrent()** returns the next member of a network group. After the call, *machinep* will contain a pointer to a string containing the name of the machine part of the network group member, and similarly for *userp* and *domainp*. If any of *machinep*, *userp* or *domainp* is returned as a NULL pointer, it signifies a wild card. **getnetgrent()** will use **malloc(3V)** to allocate space for the name. This space is released when a **endnetgrent()** call is made. **getnetgrent()** returns 1 if it succeeded in obtaining another member of the network group, 0 if it has reached the end of the group.

**getnetgrent()** establishes the network group from which **getnetgrent()** will obtain members, and also restarts calls to **getnetgrent()** from the beginning of the list. If the previous **setnetgrent()** call was to a different network group, a **endnetgrent()** call is implied. **endnetgrent()** frees the space allocated during the **getnetgrent()** calls. **inetgr** returns 1 or 0, depending on whether *netgroup* contains the machine, user, domain triple as a member. Any of the three strings *machine*, *user*, or *domain* can be NULL, in which case it signifies a wild card.

**FILES**

/etc/netgroup

**WARNINGS**

The Network Information Service (NIS) must be running when using **getnetgrent()**, since it only inspects the NIS netgroup map, never the local files.

**NOTES**

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed.

## NAME

getopt, optarg, optind – get option letter from argument vector

## SYNOPSIS

```
int getopt(argc, argv, optstring)
int argc;
char **argv;
char *optstring;

extern char *optarg;
extern int optind, opterr;
```

## DESCRIPTION

**getopt()** returns the next option letter in *argv* that matches a letter in *optstring*. *optstring* must contain the option letters the command using **getopt()** will recognize; if a letter is followed by a colon, the option is expected to have an argument, or group of arguments, which must be separated from it by white space.

*optarg* is set to point to the start of the option argument on return from **getopt**.

**getopt()** places in **optind** the *argv* index of the next argument to be processed. **optind** is external and is initialized to 1 before the first call to **getopt**.

When all options have been processed (that is, up to the first non-option argument), **getopt()** returns **-1**. The special option “—” may be used to delimit the end of the options; when it is encountered, **-1** will be returned, and “—” will be skipped.

## DIAGNOSTICS

**getopt()** prints an error message on the standard error and returns a question mark (?) when it encounters an option letter not included in *optstring* or no option-argument after an option that expects one. This error message may be disabled by setting **opterr** to 0.

## EXAMPLE

The following code fragment shows how one might process the arguments for a command that can take the mutually exclusive options **a** and **b**, and the option **o**, which requires an option argument:

```
main(argc, argv)
int argc;
char **argv;
{
    int c;
    extern char *optarg;
    extern int optind;
    .
    .
    .
    while ((c = getopt(argc, argv, "abo:")) != -1)
        switch (c) {
            case 'a':
                if (bflg)
                    errflg++;
                else
                    aflg++;
                break;
            case 'b':
                if (aflg)
                    errflg++;
                else
                    bproc ();
                break;
```

```
        case '0':
            ofile = optarg;
            break;
        case '?':
            errflg++;
    }
    if (errflg) {
        (void)fprintf(stderr, "usage: ... ");
        exit (2);
    }
    for (; optind < argc; optind++) {
        if (access(argv[optind], 4)) {
            .
            .
            .
        }
    }
```

**SEE ALSO****getopts(1)****WARNING**

Changing the value of the variable **optind**, or calling **getopt()** with different values of *argv*, may lead to unexpected results.

**NAME**

**getpass** -- read a password

**SYNOPSIS**

```
char *getpass(prompt)  
char *prompt;
```

**DESCRIPTION**

**getpass()** reads up to a NEWLINE or EOF from the file */dev/tty*, or if that cannot be opened, from the standard input, after prompting with the null-terminated string *prompt* and disabling echoing. A pointer is returned to a null-terminated string of at most 8 characters. An interrupt will terminate input and send an interrupt signal to the calling program before returning.

**SYSTEM V DESCRIPTION**

If */dev/tty* cannot be opened, **getpass()** returns a NULL pointer. It does not read the standard input.

**FILES**

*/dev/tty*

**SEE ALSO**

**crypt(3)**

**NOTES**

The above routine uses **<stdio.h>**, which increases the size of programs not otherwise using standard I/O, more than might be expected.

**BUGS**

The return value points to static data whose content is overwritten by each call.

**NAME**

getprotoent, getprotobynumber, getprotobynname, setprotoent, endprotoent – get protocol entry

**SYNOPSIS**

```
#include <netdb.h>

struct protoent *getprotoent()

struct protoent *getprotobynname(name)
char *name;

struct protoent *getprotobynumber(proto)
int proto;

setprotoent(stayopen)
int stayopen;

endprotoent()
```

**DESCRIPTION**

getprotoent, getprotobynname, and getprotobynumber() each return a pointer to an object with the following structure containing the broken-out fields of a line in the network protocol data base, */etc/protocols*.

```
struct protoent {
    char    *p_name;        /* official name of protocol */
    char    **p_aliases;    /* alias list */
    int     p_proto;        /* protocol number */
};
```

The members of this structure are:

**p\_name**                   The official name of the protocol.  
**p\_aliases**                A zero terminated list of alternate names for the protocol.  
**p\_proto**                  The protocol number.

getprotoent() reads the next line of the file, opening the file if necessary.

setprotoent() opens and rewinds the file. If the *stayopen* flag is non-zero, the net data base will not be closed after each call to getprotoent() (either directly, or indirectly through one of the other “getproto” calls).

endprotoent() closes the file.

getprotobynname() and getprotobynumber() sequentially search from the beginning of the file until a matching protocol name or protocol number is found, or until end-of-file is encountered.

**FILES**

*/etc/protocols*

**SEE ALSO**

protocols(5), ypserv(8)

**DIAGNOSTICS**

A NULL pointer is returned on end-of-file or error.

**BUGS**

All information is contained in a static area so it must be copied if it is to be saved. Only the Internet protocols are currently understood.

**NAME**

getpw – get name from uid

**SYNOPSIS**

```
getpw(uid, buf)
char *buf;
```

**DESCRIPTION**

getpw() is obsoleted by getpwent(3V).

getpw() searches the password file for the (numerical) *uid*, and fills in *buf* with the corresponding line; it returns non-zero if *uid* could not be found. The line is null-terminated.

**FILES**

/etc/passwd

**SEE ALSO**

getpwent(3V), passwd(5)

**DIAGNOSTICS**

Non-zero return on error.

**NAME**

getpwaent, getpwanam, setpwaent, endpwaent, fgetpwaent – get password adjunct file entry

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/label.h>
#include <sys/audit.h>
#include <pwdadj.h>

struct passwd_adjunct *getpwaent()

struct passwd_adjunct *getpwanam(name)
char *name;

struct passwd_adjunct *fgetpwaent(f)
FILE *f;

void setpwaent()

void endpwaent()
```

**DESCRIPTION**

Both `getpwaent()` and `getpwanam()` return a pointer to an object with the following structure containing the broken-out fields of a line in the password adjunct file. Each line in the file contains a `passwd_adjunct` structure, declared in the `<pwdadj.h>` header file:

```
struct passwd_adjunct {
    char      *pwa_name;
    char      *pwa_passwd;
    blabel_t  pwa_minimum;
    blabel_t  pwa_maximum;
    blabel_t  pwa_def;
    audit_state_t pwa_au_always;
    audit_state_t pwa_au_never;
    int       pwa_version;
};
```

When first called, `getpwaent()` returns a pointer to a `passwd_adjunct` structure describing data from the first line in the file. Thereafter, it returns a pointer to a `passwd_adjunct` structure describing data from the next line in the file. So successive calls can be used to search the entire file.

`getpwanam()` searches from the beginning of the file until it finds a login name matching *name*, then returns a pointer to the particular structure in which it was found.

Calling `setpwaent()` rewinds the password adjunct file to allow repeated searches. Calling `endpwaent()` closes the password adjunct file when processing is complete.

Because read access is required on `/etc/security/passwd.adjunct`, `getpwaent()` and `getpwanam()` will fail unless the calling process has effective UID of root.

**FILES**

```
/etc/security/passwd.adjunct
/var/yp/domainname/passwd.adjunct.byname
```

**DIAGNOSTICS**

A NULL pointer is returned on end-of-file or error.

**SEE ALSO**

`getpwent(3V)`, `getgrent(3V)`, `passwd.adjunct(5)`, `ypserv(8)`

**BUGS**

All information is contained in a static area, so it must be copied if it is to be saved.



## NAME

getpwent, getpwuid, getpwnam, setpwent, endpwent, setpwfile, fgetpwent – get password file entry

## SYNOPSIS

```
#include <pwd.h>

struct passwd *getpwent()

struct passwd *getpwuid(uid)
uid_t uid;

struct passwd *getpwnam(name)
char *name;

void setpwent()

void endpwent()

int setpwfile(name)
char *name;

struct passwd *fgetpwent(f)
FILE *f;
```

## DESCRIPTION

**getpwent()**, **getpwuid()** and **getpwnam()** each return a pointer to an object with the following structure containing the fields of a line in the password file. Each line in the file contains a **passwd** structure, declared in the **<pwd.h>** header file:

```
struct passwd {
    char    *pw_name;
    char    *pw_passwd;
    uid_t   pw_uid;
    gid_t   pw_gid;
    int     pw_quota;
    char    *pw_comment;
    char    *pw_gecos;
    char    *pw_dir;
    char    *pw_shell;
};

struct passwd *getpwent(), *getpwuid(), *getpwnam();
```

The fields **pw\_quota** and **pw\_comment** are unused; the others have meanings described in **passwd(5)**. When first called, **getpwent()** returns a pointer to the first **passwd** structure in the file; thereafter, it returns a pointer to the next **passwd** structure in the file; so successive calls can be used to search the entire file. **getpwuid()** searches from the beginning of the file until a numerical user ID matching *uid* is found and returns a pointer to the particular structure in which it was found. **getpwnam()** searches from the beginning of the file until a login name matching *name* is found, and returns a pointer to the particular structure in which it was found. If an end-of-file or an error is encountered on reading, these functions return a NULL pointer.

A call to **setpwent()** has the effect of rewinding the password file to allow repeated searches. **endpwent()** may be called to close the password file when processing is complete.

**setpwfile()** changes the default password file to *name* thus allowing alternate password files to be used. Note: it does *not* close the previous file. If this is desired, **endpwent()** should be called prior to it. **setpwfile()** will fail if it is called before a call to one of **getpwent()**, **getpwuid()**, **setpwent()**, or **getpwnam()**, or if it is called before a call to one of these functions and after a call to **endpwent()**.

**fgetpwent()** returns a pointer to the next **passwd** structure in the stream *f*, which matches the format of the password file */etc/passwd*.

**SYSTEM V DESCRIPTION**

**struct passwd** is declared in **pwd.h** as:

```
struct passwd {
    char    *pw_name;
    char    *pw_passwd;
    uid_t   pw_uid;
    gid_t   pw_gid;
    char    *pw_age;
    char    *pw_comment;
    char    *pw_gecos;
    char    *pw_dir;
    char    *pw_shell;
};
```

The field **pw\_age** is used to hold a value for “password aging” on some systems; “password aging” is not supported on Sun systems.

**RETURN VALUES**

**getpwent()**, **getpwuid()**, and **getpwnam()** return a pointer to **struct passwd** on success. On EOF or error, or if the requested entry is not found, they return **NULL**.

**setpwfile()** returns:

1        on success.  
0        on failure.

**FILES**

*/etc/passwd*  
*/var/yp/domainname/passwd.byname*  
*/var/yp/domainname/passwd.byuid*

**SEE ALSO**

**getgrent(3V)**, **issecure(3)**, **getlogin(3V)**, **passwd(5)**, **ypserv(8)**

**NOTES**

The above routines use the standard I/O library, which increases the size of programs not otherwise using standard I/O more than might be expected.

**setpwfile()** and **fgetpwent()** are obsolete and should not be used, because when the system is running in secure mode (see **issecure(3)**), the password file only contains part of the information needed for a user database entry.

**BUGS**

All information is contained in a static area which is overwritten by subsequent calls to these functions, so it must be copied if it is to be saved.

**NAME**

getrpcent, getrpcbyname, getrpcbynumber, endrpcent, setrpcent – get RPC entry

**SYNOPSIS**

```
#include <netdb.h>

struct rpcent *getrpcent()
struct rpcent *getrpcbyname(name)
char *name;
struct rpcent *getrpcbynumber(number)
int number;

setrpcent (stayopen)
int stayopen

endrpcent ()
```

**DESCRIPTION**

**getrpcent**, **getrpcbyname**, and **getrpcbynumber()** each return a pointer to an object with the following structure containing the broken-out fields of a line in the rpc program number data base, */etc/rpc*.

```
struct  rpcent {
    char   *r_name;      /* name of server for this rpc program */
    char   **r_aliases; /* alias list */
    long   r_number;    /* rpc program number */
};
```

The members of this structure are:

<b>r_name</b>	The name of the server for this rpc program.
<b>r_aliases</b>	A zero terminated list of alternate names for the rpc program.
<b>r_number</b>	The rpc program number for this service.

**getrpcent()** reads the next line of the file, opening the file if necessary.

**setrpcent()** opens and rewinds the file. If the *stayopen* flag is non-zero, the net data base will not be closed after each call to **getrpcent()** (either directly, or indirectly through one of the other “getrpc” calls).

**endrpcent** closes the file.

**getrpcbyname()** and **getrpcbynumber()** sequentially search from the beginning of the file until a matching rpc program name or program number is found, or until end-of-file is encountered.

**FILES**

*/etc/rpc*

**SEE ALSO**

**rpc(5)**, **rpcinfo(8C)**, **ypserv(8)**

**DIAGNOSTICS**

A NULL pointer is returned on EOF or error.

**BUGS**

All information is contained in a static area so it must be copied if it is to be saved.

**NAME**

gets, fgets – get a string from a stream

**SYNOPSIS**

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

```
char *gets(s)
```

```
char *s;
```

```
char *fgets(s, n, stream)
```

```
char *s;
```

```
FILE *stream;
```

**DESCRIPTION**

**gets()** reads characters from the standard input stream, **stdin**, into the array pointed to by *s*, until a NEWLINE character is read or an EOF condition is encountered. The NEWLINE character is discarded and the string is terminated with a null character. **gets()** returns its argument.

**fgets()** reads characters from the stream into the array pointed to by *s*, until *n*–1 characters are read, a NEWLINE character is read and transferred to *s*, or an EOF condition is encountered. The string is then terminated with a null character. **fgets()** returns its first argument.

**SEE ALSO**

**puts(3S)**, **getc(3V)**, **scanf(3V)**, **fread(3S)**, **ferror(3V)**

**BUGS**

If the input to **gets ()** or **fgets ()** contains a null character, the null terminates the input, and all subsequent data will be lost.

**DIAGNOSTICS**

If EOF is encountered and no characters have been read, no characters are transferred to *s* and a NULL pointer is returned. If a read error occurs, such as trying to use these functions on a file that has not been opened for reading, a NULL pointer is returned. Otherwise *s* is returned.

**NAME**

getservent, getservbyport, getservbyname, setservent, endservent – get service entry

**SYNOPSIS**

```
#include <netdb.h>

struct servent *getservent()

struct servent *getservbyname(name, proto)
char *name, *proto;

struct servent *getservbyport(port, proto)
int port; char *proto;

setservent(stayopen)
int stayopen;

endservent()
```

**DESCRIPTION**

*getservent*, *getservbyname*, and *getservbyport* each return a pointer to an object with the following structure containing the broken-out fields of a line in the network services data base, */etc/services*.

```
struct servent {
    char    *s_name;        /* official name of service */
    char    **s_aliases;    /* alias list */
    int     s_port;        /* port service resides at */
    char    *s_proto;       /* protocol to use */
};
```

The members of this structure are:

<b>s_name</b>	The official name of the service.
<b>s_aliases</b>	A zero terminated list of alternate names for the service.
<b>s_port</b>	The port number at which the service resides. Port numbers are returned in network short byte order.
<b>s_proto</b>	The name of the protocol to use when contacting the service.

*getservent()* reads the next line of the file, opening the file if necessary.

*getservent()* opens and rewinds the file. If the *stayopen* flag is non-zero, the net data base will not be closed after each call to *getservent()* (either directly, or indirectly through one of the other “*getserv*” calls).

*endservent()* closes the file.

*getservbyname()* and *getservbyport()* sequentially search from the beginning of the file until a matching protocol name or port number is found, or until end-of-file is encountered. If a protocol name is also supplied (non-NULL), searches must also match the protocol.

**FILES**

*/etc/services*

**SEE ALSO**

*getprotoent(3N)*, *services(5)*, *ypserv(8)*

**DIAGNOSTICS**

A NULL pointer is returned on end-of-file or error.

**BUGS**

All information is contained in a static area so it must be copied if it is to be saved. Expecting port numbers to fit in a 32 bit quantity is probably naive.

## NAME

getsubopt – parse sub options from a string.

## SYNOPSIS

```
int getsubopt(optionp, tokens, valuep)
char **optionp;
char *tokens[];
char **valuep;
```

## DESCRIPTION

**getsubopt()** is a function to parse suboptions in a flag argument that was initially parsed by **getopt(3)**. These suboptions are separated by commas and may consist of either a single token, or a token-value pair separated by an equal sign. Since commas delimit suboptions in the option string they are not allowed to be part of the suboption or the value of a suboption. An example command that uses this syntax is **mount(8)**, which allows you to specify mount parameters with the **-o** switch as follows :

```
pepper % mount -o rw,hard,bg,wsize=1024 speed:/usr /usr
```

In this example there are four suboptions: 'rw', 'hard', 'bg', and 'wsize', the last of which has an associated value of 1024.

**getsubopt()** takes the address of a pointer to the option string, a vector of possible tokens, and the address of a value string pointer. It returns the index of the token that matched the suboption in the input string or -1 if there was no match. If the option string at *\*optionp* contains only one suboption, **getsubopt()** updates *\*optionp* to point to the NUL at the end of the string, otherwise it isolates the suboption by replacing the comma separator with a NUL, and updates *\*optionp* to point to the start of the next suboption. If the suboption has an associated value, **getsubopt()** updates *\*valuep* to point to the value's first character. Otherwise it sets *\*valuep* to NULL.

The token vector is organized as a series of pointers to null-terminated strings. The end of the token vector is identified by a NULL pointer.

When **getsubopt()** returns, if *\*valuep* is not NULL, then the suboption processed included a value. The calling program may use this information to determine if the presence or lack of a value for this suboption is an error.

Additionally, when **getsubopt()** fails to match the suboption with the tokens in the *tokens* array, the calling program should decide if this is an error, or if the unrecognized option should be passed on to another program.

## DIAGNOSTICS

**getsubopt()** returns -1 when the token it is scanning is not in the token vector. The variable addressed by *valuep* contains a pointer to the first character of the *token* that was not recognized rather than a pointer to a value for that token.

The variable addressed by *optionp* points to the next option to be parsed, or a NUL character if there are no more options.

## EXAMPLE

The following code fragment shows how you might process options to the **mount(8)** command using **getsubopt(3)**.

```
char *myopts[] = {
#define READONLY      0
    "ro",
#define READWRITE    1
    "rw",
#define WRITESIZE    2
    "wsize",
#define READSIZE     3
    "rsize",
    NULL };
```

```

main(argc, argv)
    int argc;
    char **argv;
{
    int sc, c, errflag;
    char *options, *value;
    extern char *optarg;
    extern int optind;
    .
    .
    .
    while((c = getopt(argc, argv, "abf:o:")) != -1) {
        switch (c) {
            case 'a': /* process a option */
                break;
            case 'b': /* process b option */
                break;
            case 'f':
                ofile = optarg;
                break;
            case '?':
                errflag++;
                break;
            case 'o':
                options = optarg;
                while (*options != '\0') {
                    switch(getsubopt(&options, myopts, &value) {
                        case READONLY : /* process ro option */
                            break;
                        case READWRITE : /* process rw option */
                            break;
                        case WRITESIZE : /* process wsize option */
                            if (value == NULL) {
                                error_no_arg();
                                errflag++;
                            } else
                                write_size = atoi(value);
                            break;
                        case READSIZE : /* process rsize option */
                            if (value == NULL) {
                                error_no_arg();
                                errflag++;
                            } else
                                read_size = atoi(value);
                            break;
                        default :
                            /* process unknown token */
                            error_bad_token(value);
                            errflag++;
                            break;
                    }
                }
        }
    }
    break;
}

```

```
        }  
    }  
    if (errflag) {  
        /* print Usage instructions etc. */  
    }  
    for (; optind < argc; optind++) {  
        /* process remaining arguments */  
    }  
    .  
    .  
    .  
}
```

**SEE ALSO****getopt(3)****NOTES**

During parsing, commas in the option input string are changed to nulls.

White space in tokens or token-value pairs must be protected from the shell by quotes.



**NAME**

`gettext`, `textdomain` – retrieve a message string, get and set text domain

**SYNOPSIS**

```
char *gettext(msgtag)
char *msgtag;

char *textdomain(domainname)
char *domainname;
```

**DESCRIPTION**

`gettext()` returns a pointer to a null-terminated string (target string). *msgtag* is a string used at run-time to select the target string from the current domain of the active pool of messages. The length and contents of strings returned by `gettext()` are undetermined until called at run-time. The string returned by `gettext()` cannot be modified by the caller, but may be overwritten by a subsequent call to `gettext()`. The `LC_MESSAGES` locale category setting determines the locale of strings that `gettext()` returns.

The calling process can dynamically change the choice of locale for strings returned by `gettext()` by invoking the `setlocale(3V)` function with the correct category and the required locale. If `setlocale()` is not called or is called with an invalid value, `gettext()` defaults to the "C" locale. The default name for the current domain is the empty string.

`gettext()` first attempts to resolve the target string from the active domain and locale of the message pool. The current locale and domain are determined by the combination of both the `LC_MESSAGES` category of locale and the current domain setting.

If the target string cannot be found by using the current locale and domain then *msgtag* and current domain are applied to the implementation-defined default locale (this default locale could contain any language). If the default locale does not also contain the target string then the *msgtag* and current domain will be applied to the "C" locale of the message pool. If the target string still cannot be found then `gettext()` will return *msgtag*.

Any of the following conditions will result in a message not being found in the string archive:

- Non-existent archive selected after `setlocale()` or `textdomain()` was called.
- Non-existent archive in the "C" environment if `setlocale()` was not called.
- Non-existent or deleted entry in the archive.

`textdomain()` sets the current domain to *domainname*. Subsequent calls to `gettext()` refer to this domain. If *domainname* is NULL, `textdomain()` returns the name of the current domain without changing it.

The setting of domain made by the last successful `textdomain()` call remains valid across any number of subsequent calls to `setlocale()`.

**RETURN VALUES**

`gettext()` returns a pointer to the null-terminated target string on success. On failure, `gettext()` returns *msgtag*.

`textdomain()` returns a pointer to the name of the current domain. If the domain has not been set prior to this call, `textdomain()` returns a pointer to an empty string. `textdomain()` returns NULL if:

- *domainname* contains an invalid character.
- *domainname* is longer than `LINE_MAX` bytes in length.
- If, at the time of the call to `textdomain()`, the combination of current locale and *domainname* creates a domain that does not exist at run-time. Note: in this case `textdomain()` may have been called prior to a successful `setlocale(3V)` call, but `textdomain()` will always check against current locale setting.

**EXAMPLES**

The following produces 'Hit Return\n' in a locale that is invalid or is valid and contains the same target string as the key:

```
printf( gettext( "Hit Return\n" );
```

On a system whose default language is French, and whose process has the LC\_MESSAGES category validly set, the following might print: 'Bonjour':

```
setlocale( LC_MESSAGES, "" );  
textdomain( "Morning" );  
printf( gettext( "Welcome" );
```

If the LC\_MESSAGES category was invalidly set and the default (LC\_DEFAULT) is set to English, the last example above might print 'Good morning'. If the default is not set or is also invalid, the example would print 'Welcome'.

**SEE ALSO**

setlocale(3V), installtxt(8)

## NAME

getttyent, getttynam, setttyent, endtttyent – get ttytab file entry

## SYNOPSIS

```
#include <ttyent.h>

struct ttyent *getttyent()

struct ttyent *getttynam(name)
char *name;

setttyent()

endtttyent()
```

## DESCRIPTION

**getttyent()** and **getttynam()** each return a pointer to an object with the following structure containing the broken-out fields of a line from the tty description file.

```
struct ttyent {
    char *ty_name; /* terminal device name */
    char *ty_getty; /* command to execute, usually getty */
    char *ty_type; /* terminal type for termcap (3X) */
    int ty_status; /* status flags (see below for defines) */
    char *ty_window; /* command to start up window manager */
    char *ty_comment; /* usually the location of the terminal */
};
#define TTY_ON 0x1 /* enable logins (startup getty) */
#define TTY_SECURE 0x2 /* allow root to login */
```

**ty\_name** is the name of the character-special file in the directory **/dev**. For various reasons, it must reside in the directory **/dev**.

**ty\_getty** is the command (usually **getty(8)**) which is invoked by **init** to initialize tty line characteristics. In fact, any arbitrary command can be used; a typical use is to initiate a terminal emulator in a window system.

**ty\_type** is the name of the default terminal type connected to this tty line. This is typically a name from the **termcap(5)** data base. The environment variable **TERM** is initialized with this name by **getty(8)** or **login(1)**.

**ty\_status** is a mask of bit fields which indicate various actions to be allowed on this tty line. The following is a description of each flag.

**TTY\_ON**

Enables logins (that is, **init(8)** will start the specified “getty” command on this entry).

**TTY\_SECURE**

Allows root to login on this terminal. Note: **TTY\_ON** must be included for this to be useful.

**ty\_window** is the command to execute for a window system associated with the line. The window system will be started before the command specified in the **ty\_getty** entry is executed. If none is specified, this will be **NULL**.

**ty\_comment** is the trailing comment field, if any; a leading delimiter and white space will be removed.

**getttyent()** reads the next line from the **ttytab** file, opening the file if necessary; **setttyent()** rewinds the file; **endtttyent()** closes it.

**gettynam()** searches from the beginning of the file until a matching *name* is found (or until EOF is encountered).

**FILES**

**/etc/ttytab**

**SEE ALSO**

**login(1), ttyslot(3V), gettytab(5), ttytab(5), termcap(5), getty(8), init(8)**

**DIAGNOSTICS**

NULL pointer (0) returned on EOF or error.

**BUGS**

All information is contained in a static area so it must be copied if it is to be saved.

**NAME**

getusershell, setusershell, endusershell – get legal user shells

**SYNOPSIS**

**char \*getusershell()**

**setusershell()**

**endusershell()**

**DESCRIPTION**

**getusershell()** returns a pointer to a legal user shell as defined by the system manager in the file **/etc/shells**. If **/etc/shells** does not exist, the four locations of the two standard system shells **/bin/sh**, **/bin/csh**, **/usr/bin/sh** and **/usr/bin/csh** are returned.

**getusershell()** reads the next line (opening the file if necessary); **setusershell()** rewinds the file; **endusershell()** closes it.

**FILES**

**/etc/shells**

**/bin/sh**

**/bin/csh**

**/usr/bin/sh**

**/usr/bin/csh**

**DIAGNOSTICS**

The routine **getusershell()** returns a NULL pointer (0) on EOF or error.

**BUGS**

All information is contained in a static area so it must be copied if it is to be saved.

**NAME**

getwd – get current working directory pathname

**SYNOPSIS**

```
#include <sys/param.h>
```

```
char *getwd(pathname)  
char pathname[MAXPATHLEN];
```

**DESCRIPTION**

getwd() copies the absolute pathname of the current working directory to *pathname* and returns a pointer to the result.

**DIAGNOSTICS**

getwd() returns zero and places a message in *pathname* if an error occurs.

## NAME

hsearch, hcreate, hdestroy – manage hash search tables

## SYNOPSIS

```
#include <search.h>

ENTRY *hsearch (item, action)
ENTRY item;
ACTION action;

int hcreate (nel)
unsigned nel;

void hdestroy ( )
```

## DESCRIPTION

**hsearch()** is a hash-table search routine generalized from Knuth (6.4) Algorithm D. It returns a pointer into a hash table indicating the location at which an entry can be found. *item* is a structure of type **ENTRY** (defined in the `<search.h>` header file) containing two pointers: *item.key* points to the comparison key, and *item.data* points to any other data to be associated with that key. (Pointers to types other than character should be cast to pointer-to-character.) *action* is a member of an enumeration type **ACTION** indicating the disposition of the entry if it cannot be found in the table. **ENTER** indicates that the item should be inserted in the table at an appropriate point. **FIND** indicates that no entry should be made. Unsuccessful resolution is indicated by the return of a **NULL** pointer.

**hcreate()** allocates sufficient space for the table, and must be called before **hsearch()** is used. *nel* is an estimate of the maximum number of entries that the table will contain. This number may be adjusted upward by the algorithm in order to obtain certain mathematically favorable circumstances.

**hdestroy()** destroys the search table, and may be followed by another call to **hcreate**.

## NOTES

**hsearch()** uses **open addressing** with a *multiplicative* hash function.

## EXAMPLE

The following example will read in strings followed by two numbers and store them in a hash table, discarding duplicates. It will then read in strings and find the matching entry in the hash table and print it out.

```
#include <stdio.h>
#include <search.h>
struct info {          /* this is the info stored in the table */
    int age, room;    /* other than the key. */
};
#define
NUM_EMPL 5000 /* # of elements in search table */
main( )
{
    /* space to store strings */
    char string_space[NUM_EMPL*20];
    /* space to store employee info */
    struct info info_space[NUM_EMPL];
    /* next avail space in string_space */
    char *str_ptr = string_space;
    /* next avail space in info_space */
    struct info *info_ptr = info_space;
    ENTRY item, *found_item, *hsearch( );
    /* name to look for in table */
    char name_to_find[30];
    int i = 0;
    /* create table */
```

```

        (void) hcreate(NUM_EMPL);
        while (scanf("%s%d%d", str_ptr, &info_ptr->age,
                    &info_ptr->room) !=
EOF && i++ <
NUM_EMPL) {
            /* put info in structure, and structure in item */
            item.key = str_ptr;
            item.data = (char *)info_ptr;
            str_ptr += strlen(str_ptr) + 1;
            info_ptr++;
            /* put item into table */
            (void) hsearch(item,
ENTER);
        }
        /* access table */
        item.key = name_to_find;
        while (scanf("%s", item.key) != EOF) {
            if ((found_item = hsearch(item,
FIND)) != NULL) {
                /* if item is in the table */
                (void)printf("found %s, age = %d, room = %d\n",
                    found_item->key,
                    ((struct info *)found_item->data)->age,
                    ((struct info *)found_item->data)->room);
            } else {
                (void)printf("no such employee %s\n",
                    name_to_find);
            }
        }
    }
}

```

**SEE ALSO**

bsearch(3), lsearch(3), malloc(3V), string(3), tsearch(3)

**DIAGNOSTICS**

**hsearch()** returns a NULL pointer if either the action is **FIND** and the item could not be found or the action is **ENTER** and the table is full.

**hcreate()** returns zero if it cannot allocate sufficient space for the table.

**WARNING**

**hsearch()** and **hcreate()** use **malloc(3V)** to allocate space.

**BUGS**

Only one hash search table may be active at any given time.



**NAME**

inet inet\_addr, inet\_network, inet\_makeaddr, inet\_lnaof, inet\_netof, inet\_ntoa – Internet address manipulation

**SYNOPSIS**

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

unsigned long
inet_addr(cp)
char *cp;

inet_network(cp)
char *cp;

struct in_addr
inet_makeaddr(net, lna)
int net, lna;

inet_lnaof(in)
struct in_addr in;

inet_netof(in)
struct in_addr in;

char *
inet_ntoa(in)
struct in_addr in;
```

**DESCRIPTION**

The routines `inet_addr()` and `inet_network()` each interpret character strings representing numbers expressed in the Internet standard ‘.’ notation, returning numbers suitable for use as Internet addresses and Internet network numbers, respectively. The routine `inet_makeaddr()` takes an Internet network number and a local network address and constructs an Internet address from it. The routines `inet_netof()` and `inet_lnaof()` break apart Internet host addresses, returning the network number and local network address part, respectively.

The routine `inet_ntoa()` returns a pointer to a string in the base 256 notation “d.d.d.d” described below.

All Internet address are returned in network order (bytes ordered from left to right). All network numbers and local address parts are returned as machine format integer values.

**INTERNET ADDRESSES**

Values specified using the ‘.’ notation take one of the following forms:

```
a.b.c.d
a.b.c
a.b
a
```

When four parts are specified, each is interpreted as a byte of data and assigned, from left to right, to the four bytes of an Internet address. Note: when an Internet address is viewed as a 32-bit integer quantity on Sun386i systems, the bytes referred to above appear as **d.c.b.a**. That is, Sun386i bytes are ordered from right to left.

When a three part address is specified, the last part is interpreted as a 16-bit quantity and placed in the right most two bytes of the network address. This makes the three part address format convenient for specifying Class B network addresses as “128.net.host”.

When a two part address is supplied, the last part is interpreted as a 24-bit quantity and placed in the right most three bytes of the network address. This makes the two part address format convenient for specifying Class A network addresses as "net.host".

When only one part is given, the value is stored directly in the network address without any byte rearrangement.

All numbers supplied as "parts" in a '.' notation may be decimal, octal, or hexadecimal, as specified in the C language (that is, a leading 0x or 0X implies hexadecimal; otherwise, a leading 0 implies octal; otherwise, the number is interpreted as decimal).

**SEE ALSO**

**gethostent(3N), getnetent(3N), hosts(5), networks(5),**

**DIAGNOSTICS**

The value -1 is returned by **inet\_addr()** and **inet\_network()** for malformed requests.

**BUGS**

The problem of host byte ordering versus network byte ordering is confusing. A simple way to specify Class C network addresses in a manner similar to that for Class B and Class A is needed.

The return value from **inet\_ntoa()** points to static information which is overwritten in each call.

**NAME**

**initgroups** – initialize supplementary group IDs

**SYNOPSIS**

```
initgroups(name, basegid)  
char *name;  
int basegid;
```

**DESCRIPTION**

**initgroups()** reads through the group file and sets up, using the **setgroups** call (see **getgroups(2V)**), the supplementary group IDs for the user specified in *name*. The **basegid** is automatically included in the supplementary group IDs. Typically this value is given as the group number from the password file.

**FILES**

**/etc/group**

**SEE ALSO**

**getgroups(2V)**, **getgrent(3V)**

**DIAGNOSTICS**

**initgroups()** returns **-1** if it was not invoked by the super-user.

**BUGS**

**initgroups()** uses the routines based on **getgrent(3V)**. If the invoking program uses any of these routines, the group structure will be overwritten in the call to **initgroups**.

**NAME**

**insque, remque** – insert/remove element from a queue

**SYNOPSIS**

```
struct qelem {
    struct qelem *q_forw;
    struct qelem *q_back;
    char    q_data[];
};
```

```
insque(elem, pred)
struct qelem *elem, *pred;
```

```
remque(elem)
struct qelem *elem;
```

**DESCRIPTION**

**insque()** and **remque()** manipulate queues built from doubly linked lists. Each element in the queue must be in the form of "struct qelem". **insque()** inserts *elem* in a queue immediately after *pred*; **remque()** removes an entry *elem* from a queue.

**NAME**

issecure – indicates whether system is running secure

**SYNOPSIS**

```
int issecure()
```

**DESCRIPTION**

This function tells whether the system has been configured to run in secure mode. It returns 0 if the system is not running secure, and non-zero if the system is running secure.

**NAME**

`kvm_getu`, `kvm_getcmd` – get the u-area or invocation arguments for a process

**SYNOPSIS**

```
#include <kvm.h>
#include <sys/param.h>
#include <sys/user.h>
#include <sys/proc.h>

struct user *kvm_getu(kd, proc)
kvm_t *kd;
struct proc *proc;

int kvm_getcmd(kd, proc, u, arg, env)
kvm_t *kd;
struct proc *proc;
struct user *u;
char ***arg;
char ***env;
```

**DESCRIPTION**

`kvm_getu()` reads the u-area of the process specified by *proc* to an area of static storage associated with *kd* and returns a pointer to it. Subsequent calls to `kvm_getu()` will overwrite this static area.

*kd* is a pointer to a kernel identifier returned by `kvm_open(3K)`. *proc* is a pointer to a copy (in the current process' address space) of a *proc* structure (obtained, for instance, by a prior `kvm_nextproc(3K)` call).

`kvm_getcmd()` constructs a list of string pointers that represent the command arguments and environment that were used to initiate the process specified by *proc*.

*kd* is a pointer to a kernel identifier returned by `kvm_open(3K)`. *u* is a pointer to a copy (in the current process' address space) of a *user* structure (obtained, for instance, by a prior `kvm_getu()` call). If *arg* is not NULL, then the command line arguments are formed into a null-terminated array of string pointers. The address of the first such pointer is returned in *arg*. If *env* is not NULL, then the environment is formed into a null-terminated array of string pointers. The address of the first of these is returned in *env*.

The pointers returned in *arg* and *env* refer to data allocated by `malloc(3V)` and should be freed (by a call to `free` (see `malloc(3V)`) when no longer needed. Both the string pointers and the strings themselves are deallocated when freed.

Since the environment and command line arguments may have been modified by the user process, there is no guarantee that it will be possible to reconstruct the original command at all. Thus, `kvm_getcmd()` will make the best attempt possible, returning `-1` if the user process data is unrecognizable.

**RETURN VALUES**

On success, `kvm_getu()` returns a pointer to a copy of the u-area of the process specified by *proc*. On failure, it returns NULL.

`kvm_getcmd()` returns:

```
0      on success.
-1     on failure.
```

**SEE ALSO**

`execve(2V)`, `kvm_nextproc(3K)`, `kvm_open(3K)`, `kvm_read(3K)`, `malloc(3V)`

**NOTES**

If `kvm_getcmd()` returns `-1`, the caller still has the option of using the command line fragment that is stored in the `u-area`.

**NAME**

kvm\_getproc, kvm\_nextproc, kvm\_setproc – read system process structures

**SYNOPSIS**

```
#include <kvm.h>
#include <sys/param.h>
#include <sys/time.h>
#include <sys/proc.h>

struct proc *kvm_getproc(kd, pid)
kvm_t *kd;
int pid;

struct proc *kvm_nextproc(kd)
kvm_t *kd;

int kvm_setproc(kd)
kvm_t *kd;
```

**DESCRIPTION**

**kvm\_nextproc()** may be used to sequentially read all of the system process structures from the kernel identified by *kd* (see **kvm\_open(3K)**). Each call to **kvm\_nextproc()** returns a pointer to the static memory area that contains a copy of the next valid process table entry. There is no guarantee that the data will remain valid across calls to **kvm\_nextproc()**, **kvm\_setproc()**, or **kvm\_getproc()**. Therefore, if the process structure must be saved, it should be copied to non-volatile storage.

For performance reasons, many implementations will cache a set of system process structures. Since the system state is liable to change between calls to **kvm\_nextproc()**, and since the cache may contain obsolete information, there is no guarantee that *every* process structure returned refers to an active process, nor is it certain that *all* processes will be reported.

**kvm\_setproc()** rewinds the process list, enabling **kvm\_nextproc()** to rescan from the beginning of the system process table. **kvm\_setproc()** will always flush the process structure cache, allowing an application to re-scan the process table of a running system.

**kvm\_getproc()** locates the **proc** structure of the process specified by *pid* and returns a pointer to it. **kvm\_getproc()** does not interact with the process table pointer manipulated by **kvm\_nextproc**, however, the restrictions regarding the validity of the data still apply.

**RETURN VALUES**

On success, **kvm\_nextproc()** returns a pointer to a copy of the next valid process table entry. On failure, it returns NULL.

On success, **kvm\_getproc()** returns a pointer to the **proc** structure of the process specified by *pid*. On failure, it returns NULL.

**kvm\_setproc()** returns:

```
0      on success.
-1     on failure.
```

**SEE ALSO**

**kvm\_getu(3K)**, **kvm\_open(3K)**, **kvm\_read(3K)**



**NAME**

`kvm_nlist` – get entries from kernel symbol table

**SYNOPSIS**

```
#include <kvm.h>
#include <nlist.h>

int kvm_nlist(kd, nl)
kvm_t *kd;
struct nlist *nl;
```

**DESCRIPTION**

`kvm_nlist()` examines the symbol table from the kernel image identified by *kd* (see `kvm_open(3K)`) and selectively extracts a list of values and puts them in the array of `nlist()` structures pointed to by *nl*. The name list pointed to by `nl()` consists of an array of structures containing names, types and values. The *n\_name* field of each such structure is taken to be a pointer to a character string representing a symbol name. The list is terminated by an entry with a NULL pointer (or a pointer to a null string) in the *n\_name* field. For each entry in *nl*, if the named symbol is present in the kernel symbol table, its value and type are placed in the *n\_value* and *n\_type* fields. If a symbol cannot be located, the corresponding *n\_type* field of `nl()` is set to zero.

**RETURN VALUES**

On success, `kvm_nlist()` returns the number of symbols that were not located in the symbol table. On failure, it returns `-1` and sets all of the *n\_type* fields in members of the array pointed to by *nl* to zero.

**SEE ALSO**

`kvm_open(3K)`, `kvm_read(3K)`, `nlist(3V)`, `a.out(5)`

## NAME

`kvm_open`, `kvm_close` – specify a kernel to examine

## SYNOPSIS

```
#include <kvm.h>
#include <fcntl.h>

kvm_t *kvm_open(namelist, corefile, swapfile, flag, errstr)
char *namelist, *corefile, *swapfile;
int flag;
char *errstr;

int kvm_close(kd)
kvm_t *kd;
```

## DESCRIPTION

`kvm_open()` initializes a set of file descriptors to be used in subsequent calls to kernel VM routines. It returns a pointer to a kernel identifier that must be used as the *kd* argument in subsequent kernel VM function calls.

The *namelist* argument specifies an unstripped executable file whose symbol table will be used to locate various offsets in *corefile*. If *namelist* is NULL, the symbol table of the currently running kernel is used to determine offsets in the core image. In this case, it is up to the implementation to select an appropriate way to resolve symbolic references (for instance, using `/vmunix` as a default *namelist* file).

*corefile* specifies a file that contains an image of physical memory, for instance, a kernel crash dump file (see `savecore(8)`) or the special device `/dev/mem`. If *corefile* is NULL, the currently running kernel is accessed (using `/dev/mem` and `/dev/kmem`).

*swapfile* specifies a file that represents the swap device. If both *corefile* and *swapfile* are NULL, the swap device of the “currently running kernel” is accessed. Otherwise, if *swapfile* is NULL, `kvm_open()` may succeed but subsequent `kvm_getu(3K)` function calls may fail if the desired information is swapped out.

*flag* is used to specify read or write access for *corefile* and may have one of the following values:

<code>O_RDONLY</code>	open for reading
<code>O_RDWR</code>	open for reading and writing

*errstr* is used to control error reporting. If it is a NULL pointer, no error messages will be printed. If it is non-NULL, it is assumed to be the address of a string that will be used to prefix error messages generated by `kvm_open`. Errors are printed to `stderr`. A useful value to supply for *errstr* would be `argv[0]`. This has the effect of printing the process name in front of any error messages.

`kvm_close()` closes all file descriptors that were associated with *kd*. These files are also closed on `exit(2v)` and `execve(2V)`. `kvm_close()` also resets the `proc` pointer associated with `kvm_nextproc(3K)` and flushes any cached kernel data.

## RETURN VALUES

`kvm_open()` returns a non-NULL value suitable for use with subsequent kernel VM function calls. On failure, it returns NULL and no files are opened.

`kvm_close()` returns:

0	on success.
-1	on failure.

**FILES**

**/vmunix**  
**/dev/kmem**  
**/dev/mem**  
**/dev/drum**

**SEE ALSO**

**execve(2V), exit(2v), kvm\_getu(3K), kvm\_nextproc(3K), kvm\_nlist(3K), kvm\_read(3K), savecore(8)**

**NAME**

`kvm_read`, `kvm_write` – copy data to or from a kernel image or running system

**SYNOPSIS**

```
#include <kvm.h>

int kvm_read(kd, addr, buf, nbytes)
kvm_t *kd;
unsigned long addr;
char *buf;
unsigned nbytes;

int kvm_write(kd, addr, buf, nbytes)
kvm_t *kd;
unsigned long addr;
char *buf;
unsigned nbytes;
```

**DESCRIPTION**

`kvm_read()` transfers data from the kernel image specified by `kd` (see `kvm_open(3K)`) to the address space of the process. `nbytes` bytes of data are copied from the kernel virtual address given by `addr` to the buffer pointed to by `buf`.

`kvm_write()` is like `kvm_read()`, except that the direction of data transfer is reversed. In order to use this function, the `kvm_open(3K)` call that returned `kd` must have specified write access. If a user virtual address is given, it is resolved in the address space of the process specified in the most recent `kvm_getu(3K)` call.

**RETURN VALUES**

On success, `kvm_read()` and `kvm_write()` return the number of bytes actually transferred. On failure, they return `-1`.

**SEE ALSO**

`kvm_getu(3K)`, `kvm_nlist(3K)`, `kvm_open(3K)`

**NAME**

**l3tol, ltol3** – convert between 3-byte integers and long integers

**SYNOPSIS**

```
#include <stdlib.h>  
void l3tol (lp, cp, n)  
long *lp;  
const char *cp;  
int n;  
  
void ltol3 (cp, lp, n)  
char *cp;  
const long *lp;  
int n;
```

**DESCRIPTION**

**l3tol()** converts a list of *n* three-byte integers packed into a character string pointed to by *cp* into a list of long integers pointed to by *lp*.

**ltol3()** performs the reverse conversion from long integers (*lp*) to three-byte integers (*cp*).

These functions are useful for filesystem maintenance where the block numbers are three bytes long.

**SEE ALSO**

fs(5)

**WARNINGS**

Because of possible differences in byte ordering, the numerical values of the long integers are machine-dependent.

**NAME**

**ldahread** – read the archive header of a member of a COFF archive file

**SYNOPSIS**

```
#include <stdio.h>
#include <ar.h>
#include <filehdr.h>
#include <ldfcn.h>

int ldahread (ldptr, arhead)
LDFILE *ldptr;
ARCHDR *arhead;
```

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

If **TYPE(ldptr)** is the archive file magic number, **ldahread** reads the archive header of the COFF file currently associated with *ldptr* into the area of memory beginning at *arhead*.

**ldahread** returns **SUCCESS** or **FAILURE**. **ldahread** will fail if **TYPE(ldptr)** does not represent an archive file, or if it cannot read the archive header.

The program must be loaded with the object file access routine library **libld.a**.

**SEE ALSO**

**ldclose(3X)**, **ldfcn(3)**, **ldopen(3X)**, **intro(5)**

**NAME**

`ldclose`, `ldaclose` – close a COFF file

**SYNOPSIS**

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

int ldclose (ldptr)
LDFILE *ldptr;

int ldaclose (ldptr)
LDFILE *ldptr;
```

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

`ldopen(3X)` and `ldclose()` are designed to provide uniform access to both simple COFF object files and COFF object files that are members of archive files. Thus an archive of COFF files can be processed as if it were a series of simple COFF files.

If `TYPE(ldptr)` does not represent an archive file, `ldclose()` will close the file and free the memory allocated to the `LDFILE` structure associated with `ldptr`. If `TYPE(ldptr)` is the magic number of an archive file, and if there are any more files in the archive, `ldclose()` will reinitialize `OFFSET(ldptr)` to the file address of the next archive member and return `FAILURE`. The `LDFILE` structure is prepared for a subsequent `ldopen(3X)`. In all other cases, `ldclose()` returns `SUCCESS`.

`ldaclose()` closes the file and frees the memory allocated to the `LDFILE` structure associated with `ldptr` regardless of the value of `TYPE(ldptr)`. `ldaclose()` always returns `SUCCESS`. The function is often used in conjunction with `ldaopen`.

The program must be loaded with the object file access routine library `libld.a`.

`intro(5)` describes `INCDIR` and `LIBDIR`.

**SEE ALSO**

`fclose(3V)`, `ldfcn(3)`, `ldopen(3X)`, `intro(5)`

**NAME**

ldfcn – common object file access routines

**SYNOPSIS**

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

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

These routines are for reading COFF object files and archives containing COFF object files. Although the calling program must know the detailed structure of the parts of the object file that it processes, the routines effectively insulate the calling program from knowledge of the overall structure of the object file.

The interface between the calling program and the object file access routines is based on the defined type **LDFILE**, defined as **struct ldfile**, declared in the header file **ldfcn.h**. The primary purpose of this structure is to provide uniform access to both simple object files and to object files that are members of an archive file.

The function **ldopen(3X)** allocates and initializes the **LDFILE** structure and returns a pointer to the structure to the calling program. The fields of the **LDFILE** structure may be accessed individually through macros defined in **ldfcn.h** and contain the following information:

<b>LDFILE</b>	<b>*ldptr;</b>
<b>TYPE(ldptr)</b>	The file magic number used to distinguish between archive members and simple object files.
<b>IOPTR(ldptr)</b>	The file pointer returned by <i>fopen</i> and used by the standard input/output functions.
<b>OFFSET(ldptr)</b>	The file address of the beginning of the object file; the offset is non-zero if the object file is a member of an archive file.
<b>HEADER(ldptr)</b>	The file header structure of the object file.

The object file access functions themselves may be divided into four categories:

## (1) Functions that open or close an object file

**ldopen(3X)** and **ldaopen()** (see **ldopen(3X)**)  
 open a common object file  
**ldclose(3X)** and **ldaclose()** (see **ldclose(3X)**)  
 close a common object file

## (2) Functions that read header or symbol table information

**ldahread(3X)**  
 read the archive header of a member of an archive file  
**ldfhread(3X)**  
 read the file header of a common object file  
**ldshread(3X)** and **ldnshread()** (see **ldshread(3X)**)  
 read a section header of a common object file  
**ldtbread(3X)**  
 read a symbol table entry of a common object file  
**ldgetname(3X)**  
 retrieve a symbol name from a symbol table entry or from the string table



(3) Functions that position an object file at (seek to) the start of the section, relocation, or line number information for a particular section.

**ldohseek(3X)**

seek to the optional file header of a common object file

**ldsseek(3X) and ldnsseek()** (see **ldsseek(3X)**)

seek to a section of a common object file

**ldrseek(3X) and ldnrseek()** (see **ldrseek(3X)**)

seek to the relocation information for a section of a common object file

**ldlseek(3X) and ldlnseek()** (see **ldlseek(3X)**)

seek to the line number information for a section of a common object file

**ldtbseek(3X)**

seek to the symbol table of a common object file

(4) The unctio**n ldtbindex(3X)**, which returns the index of a particular common object file symbol table entry.

These functions are described in detail on their respective manual pages.

All the functions except **ldopen(3X)**, **ldgetname(3X)**, **ldtbindex(3X)** return either **SUCCESS** or **FAILURE**, both constants defined in **ldfcn.h**. **ldopen(3X)** and **ldaopen()** (see **ldopen(3X)**) both return pointers to an **LDFILE** structure.

Additional access to an object file is provided through a set of macros defined in **ldfcn.h**. These macros parallel the standard input/output file reading and manipulating functions, translating a reference of the **LDFILE** structure into a reference to its file descriptor field.

The following macros are provided:

**GETC(ldptr)**

**FGETC(ldptr)**

**GETW(ldptr)**

**UNGETC(c, ldptr)**

**FGETS(s, n, ldptr)**

**FREAD((char \*) ptr, sizeof (\*ptr), nitens, ldptr)**

**FSEEK(ldptr, offset, ptrname)**

**FTELL(ldptr)**

**REWIND(ldptr)**

**FEOF(ldptr)**

**FERROR(ldptr)**

**FILENO(ldptr)**

**SETBUF(ldptr, buf)**

**STROFFSET(ldptr)**

The **STROFFSET** macro calculates the address of the string table. See the manual entries for the corresponding standard input/output library functions for details on the use of the rest of the macros.

The program must be loaded with the object file access routine library **libld.a**.

**SEE ALSO**

**fseek(3S)**, **ldahread(3X)**, **ldclose(3X)**, **ldgetname(3X)**, **ldhread(3X)**, **ldhread(3X)**, **ldlseek(3X)**, **ldohseek(3X)**, **ldopen(3X)**, **ldrseek(3X)**, **ldlseek(3X)**, **ldhread(3X)**, **ldtbindex(3X)**, **ldtbread(3X)**, **ldtbseek(3X)**, **stdio(3V)**, **intro(5)**

**WARNING**

The macro **FSEEK** defined in the header file **ldfcn.h** translates into a call to the standard input/output function **fseek(3S)**. **FSEEK** should not be used to seek from the end of an archive file since the end of an archive file may not be the same as the end of one of its object file members.

**NAME**

`ldfhread` – read the file header of a COFF file

**SYNOPSIS**

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

int ldfhread (ldptr, filehead)
LDFILE *ldptr;
FILHDR *filehead;
```

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

`ldfhread()` reads the file header of the COFF file currently associated with *ldptr* into the area of memory beginning at *filehead*.

`ldfhread()` returns **SUCCESS** or **FAILURE**. `ldfhread()` will fail if it cannot read the file header.

In most cases the use of `ldfhread()` can be avoided by using the macro `HEADER(ldptr)` defined in `ldfcn.h` (see `ldfcn(3)`). The information in any field, *fieldname*, of the file header may be accessed using `HEADER(ldptr).fieldname`.

The program must be loaded with the object file access routine library `libld.a`.

**SEE ALSO**

`ldclose(3X)`, `ldfcn(3)`, `ldopen(3X)`

**NAME**

**ldgetname** – retrieve symbol name for COFF file symbol table entry

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <syms.h>
#include <ldfcn.h>

char *ldgetname (ldptr, symbol)
LDFILE *ldptr;
SYMENT *symbol;
```

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

**ldgetname()** returns a pointer to the name associated with **symbol** as a string. The string is contained in a static buffer local to **ldgetname()** that is overwritten by each call to **ldgetname()**, and therefore must be copied by the caller if the name is to be saved.

**ldgetname()** can be used to retrieve names from object files without any backward compatibility problems. **ldgetname()** will return NULL (defined in **stdio.h**) for an object file if the name cannot be retrieved. This situation can occur:

- if the “string table” cannot be found,
- if not enough memory can be allocated for the string table,
- if the string table appears not to be a string table (for example, if an auxiliary entry is handed to **ldgetname()** that looks like a reference to a name in a nonexistent string table), or
- if the name’s offset into the string table is past the end of the string table.

Typically, **ldgetname()** will be called immediately after a successful call to **ldtbread()** to retrieve the name associated with the symbol table entry filled by **ldtbread()**.

The program must be loaded with the object file access routine library **libld.a**.

**SEE ALSO**

**ldclose(3X)**, **ldfcn(3)**, **ldopen(3X)**, **ldtbread(3X)**, **ldtbseek(3X)**

## NAME

`ldlread`, `ldlinit`, `ldlitem` – manipulate line number entries of a COFF file function

## SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <linenum.h>
#include <ldfcn.h>

int ldlread(ldptr, fcnindx, linenum, linent)
LDFILE *ldptr;
long fcnindx;
unsigned short linenum;
LINENO *linent;

int ldlinit(ldptr, fcnindx)
LDFILE *ldptr;
long fcnindx;

int ldlitem(ldptr, linenum, linent)
LDFILE *ldptr;
unsigned short linenum;
LINENO *linent;
```

## AVAILABILITY

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

## DESCRIPTION

`ldlread()` searches the line number entries of the COFF file currently associated with `ldptr`. `ldlread()` begins its search with the line number entry for the beginning of a function and confines its search to the line numbers associated with a single function. The function is identified by `fcnindx`, the index of its entry in the object file symbol table. `ldlread()` reads the entry with the smallest line number equal to or greater than `linenum` into the memory beginning at `linent`.

`ldlinit()` and `ldlitem()` together perform exactly the same function as `ldlread()`. After an initial call to `ldlread()` or `ldlinit()`, `ldlitem()` may be used to retrieve a series of line number entries associated with a single function. `ldlinit()` simply locates the line number entries for the function identified by `fcnindx`. `ldlitem()` finds and reads the entry with the smallest line number equal to or greater than `linenum` into the memory beginning at `linent()`.

`ldlread()`, `ldlinit()`, and `ldlitem()` each return either SUCCESS or FAILURE. `ldlread()` will fail if there are no line number entries in the object file, if `fcnindx` does not index a function entry in the symbol table, or if it finds no line number equal to or greater than `linenum`. `ldlinit()` will fail if there are no line number entries in the object file or if `fcnindx` does not index a function entry in the symbol table. `ldlitem()` will fail if it finds no line number equal to or greater than `linenum`.

The programs must be loaded with the object file access routine library `libld.a`.

## SEE ALSO

`ldclose(3X)`, `ldfcn(3)`, `ldopen(3X)`, `ldtbindex(3X)`

**NAME**

**ldlseek, ldnlseek** – seek to line number entries of a section of a COFF file

**SYNOPSIS**

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

int ldlseek (ldptr, sectindx)
LDFILE *ldptr;
unsigned short sectindx;

int ldnlseek (ldptr, sectname)
LDFILE *ldptr;
char *sectname;
```

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

**ldlseek()** seeks to the line number entries of the section specified by *sectindx* of the COFF file currently associated with *ldptr*.

**ldnlseek()** seeks to the line number entries of the section specified by *sectname*.

**ldlseek()** and **ldnlseek()** return **SUCCESS** or **FAILURE**. **ldlseek()** will fail if *sectindx* is greater than the number of sections in the object file; **ldnlseek()** will fail if there is no section name corresponding with *\*sectname*. Either function will fail if the specified section has no line number entries or if it cannot seek to the specified line number entries.

Note that the first section has an index of **one**.

The program must be loaded with the object file access routine library **libld.a**.

**SEE ALSO**

**ldclose(3X), ldfcn(3), ldopen(3X), ldshread(3X)**

**NAME**

`ldohseek` – seek to the optional file header of a COFF file

**SYNOPSIS**

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

int ldohseek (ldptr)
LDFILE *ldptr;
```

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

`ldohseek()` seeks to the optional file header of the COFF file currently associated with *ldptr*.

`ldohsee()` returns **SUCCESS** or **FAILURE**. `ldohseek()` will fail if the object file has no optional header or if it cannot seek to the optional header.

The program must be loaded with the object file access routine library `libld.a`.

**SEE ALSO**

`ldclose(3X)`, `ldfcn(3)`, `ldopen(3X)`, `ldhread(3X)`

## NAME

`ldopen`, `ldaopen` – open a COFF file for reading

## SYNOPSIS

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

LDFILE *ldopen (filename, ldptr)
char *filename;
LDFILE *ldptr;

LDFILE *ldaopen (filename, oldptr)
char *filename;
LDFILE *oldptr;
```

## AVAILABILITY

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

## DESCRIPTION

`ldopen()` and `ldclose(3X)` are designed to provide uniform access to both simple object files and object files that are members of archive files. Thus an archive of COFF files can be processed as if it were a series of simple COFF files.

If `ldptr` has the value `NULL`, then `ldopen()` will open `filename` and allocate and initialize the `LDFILE` structure, and return a pointer to the structure to the calling program.

If `ldptr` is valid and if `TYPE(ldptr)` is the archive magic number, `ldopen()` will reinitialize the `LDFILE` structure for the next archive member of `filename`.

`ldopen()` and `ldclose(3X)` are designed to work in concert. `ldclose` will return `FAILURE` only when `TYPE(ldptr)` is the archive magic number and there is another file in the archive to be processed. Only then should `ldopen()` be called with the current value of `ldptr`. In all other cases, in particular whenever a new `filename` is opened, `ldopen()` should be called with a `NULL` `ldptr` argument.

The following is a prototype for the use of `ldopen()` and `ldclose(3X)`.

```
/* for each filename to be processed */
ldptr = NULL;
do
{
    if ( (ldptr = ldopen(filename, ldptr)) != NULL )
    {
        /* check magic number */
        /* process the file */
    }
} while (ldclose(ldptr) == FAILURE );
```

If the value of `oldptr` is not `NULL`, `ldaopen()` will open `filename` anew and allocate and initialize a new `LDFILE` structure, copying the `TYPE`, `OFFSET`, and `HEADER` fields from `oldptr`. `ldaopen()` returns a pointer to the new `LDFILE` structure. This new pointer is independent of the old pointer, `oldptr`. The two pointers may be used concurrently to read separate parts of the object file. For example, one pointer may be used to step sequentially through the relocation information, while the other is used to read indexed symbol table entries.

Both **ldopen()** and **ldaopen()** open *filename* for reading. Both functions return NULL if *filename* cannot be opened, or if memory for the **LDFILE** structure cannot be allocated. A successful open does not insure that the given file is a COFF file or an archived object file.

The program must be loaded with the object file access routine library **libld.a**.

**SEE ALSO**

**fopen(3V), ldclose(3X), ldfcn(3)**



**NAME**

`ldrseek`, `ldnrseek` – seek to relocation entries of a section of a COFF file

**SYNOPSIS**

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

int ldrseek (ldptr, sectindx)
LDFILE *ldptr;
unsigned short sectindx;

int ldnrseek (ldptr, sectname)
LDFILE *ldptr;
char *sectname;
```

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

`ldrseek()` seeks to the relocation entries of the section specified by *sectindx* of the COFF file currently associated with *ldptr*.

`ldnrseek()` seeks to the relocation entries of the section specified by *sectname*.

`ldrseek()` and `ldnrseek()` return **SUCCESS** or **FAILURE**. `ldrseek()` will fail if *sectindx* is greater than the number of sections in the object file; `ldnrseek()` will fail if there is no section name corresponding with *sectname*. Either function will fail if the specified section has no relocation entries or if it cannot seek to the specified relocation entries.

Note: the first section has an index of **one**.

The program must be loaded with the object file access routine library `libld.a`.

**SEE ALSO**

`ldclose(3X)`, `ldfcn(3)`, `ldopen(3X)`, `ldshread(3X)`

**NAME**

ldshread, ldnsbread – read an indexed/named section header of a COFF file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <scnhdr.h>
#include <ldfcn.h>

int ldshread (ldptr, sectindx, secthead)
LDFILE *ldptr;
unsigned short sectindx;
SCNHDR *secthead;

int ldnsbread (ldptr, sectname, secthead)
LDFILE *ldptr;
char *sectname;
SCNHDR *secthead;
```

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

**ldshread()** reads the section header specified by *sectindx* of the COFF file currently associated with *ldptr* into the area of memory beginning at *secthead*.

**ldnsbread()** reads the section header specified by *sectname* into the area of memory beginning at *secthead*.

**ldshread()** and **ldnsbread()** return SUCCESS or FAILURE. **ldshread()** will fail if *sectindx* is greater than the number of sections in the object file; **ldnsbread()** will fail if there is no section name corresponding with *sectname*. Either function will fail if it cannot read the specified section header.

Note: the first section header has an index of *one*.

The program must be loaded with the object file access routine library **libld.a**.

**SEE ALSO**

**ldclose(3X)**, **ldfcn(3)**, **ldopen(3X)**

**NAME**

`ldsseek`, `ldnsseek` – seek to an indexed/named section of a COFF file

**SYNOPSIS**

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

int ldsseek (ldptr, sectindx)
LDFILE *ldptr;
unsigned short sectindx;

int ldnsseek (ldptr, sectname)
LDFILE *ldptr;
char *sectname;
```

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

`ldsseek()` seeks to the section specified by *sectindx* of the COFF file currently associated with *ldptr*.

`ldnsseek()` seeks to the section specified by *sectname*.

`ldsseek()` and `ldnsseek()` return **SUCCESS** or **FAILURE**. `ldsseek()` will fail if *sectindx* is greater than the number of sections in the object file; `ldnsseek()` will fail if there is no section name corresponding with *sectname*. Either function will fail if there is no section data for the specified section or if it cannot seek to the specified section.

Note: the first section has an index of *one*.

The program must be loaded with the object file access routine library **libld.a**.

**SEE ALSO**

`ldclose(3X)`, `ldfcn(3)`, `ldopen(3X)`, `ldshread(3X)`

**NAME**

`ldtbindex` – compute the index of a symbol table entry of a COFF file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <syms.h>
#include <ldfcn.h>

long ldtbindex (ldptr)
LDFILE *ldptr;
```

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

`ldtbindex()` returns the (**long**) index of the symbol table entry at the current position of the COFF file associated with `ldptr`.

The index returned by `ldtbindex()` may be used in subsequent calls to `ldtbread(3X)`. However, since `ldtbindex()` returns the index of the symbol table entry that begins at the current position of the object file, if `ldtbindex()` is called immediately after a particular symbol table entry has been read, it will return the index of the next entry.

`ldtbindex()` will fail if there are no symbols in the object file, or if the object file is not positioned at the beginning of a symbol table entry.

Note that the first symbol in the symbol table has an index of *zero*.

The program must be loaded with the object file access routine library `libld.a`.

**SEE ALSO**

`ldclose(3X)`, `ldfcn(3)`, `ldopen(3X)`, `ldtbread(3X)`, `ldtbseek(3X)`

**NAME**

ldtbread – read an indexed symbol table entry of a COFF file

**SYNOPSIS**

```
#include <stdio.h>
#include <filehdr.h>
#include <syms.h>
#include <ldfcn.h>

int ldtbread (ldptr, symindex, symbol)
LDFILE *ldptr;
long symindex;
SYMENT *symbol;
```

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

**ldtbread()** reads the symbol table entry specified by *symindex* of the COFF file currently associated with *ldptr* into the area of memory beginning at *symbol*.

**ldtbread()** returns SUCCESS or FAILURE. **ldtbread()** will fail if *symindex* is greater than or equal to the number of symbols in the object file, or if it cannot read the specified symbol table entry.

Note: the first symbol in the symbol table has an index of *zero*.

The program must be loaded with the object file access routine library **libld.a**.

**SEE ALSO**

**ldclose(3X)**, **ldfcn(3)**, **ldopen(3X)**, **ldtbseek(3X)**, **ldgetname(3X)**

**NAME**

`ldtbseek` – seek to the symbol table of a COFF file

**SYNOPSIS**

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

int ldtbseek (ldptr)
LDFILE *ldptr;
```

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

`ldtbseek()` seeks to the symbol table of the COFF file currently associated with `ldptr`.

`ldtbseek()` returns SUCCESS or FAILURE. `ldtbseek()` will fail if the symbol table has been stripped from the object file, or if it cannot seek to the symbol table.

The program must be loaded with the object file access routine library `libld.a`.

**SEE ALSO**

`ldclose(3X)`, `ldfcn(3)`, `ldopen(3X)`, `ldtbread(3X)`

**NAME**

`localdtconv` – get date and time formatting conventions

**SYNOPSIS**

```
#include <locale.h>

struct dtconv *localdtconv()
```

**DESCRIPTION**

`localdtconv()` returns a pointer to a structure of type `struct dtconv` containing values appropriate for the formatting of dates and times according to the rules of the current locale.

The members include the following:

**char \*abbrev\_month\_names[12]**

The abbreviated names of the months; for example, the abbreviated name for January is `abbrev_month_names[0]` and the abbreviated name for December is `abbrev_month_names[11]`.

**char \*month\_names[12]**

The full names of the months; for example, the full name for January is `month_names[0]` and the full name for December is `month_names[11]`.

**char \*abbrev\_weekday\_names[7]**

The abbreviated names of the weekdays; for example, the abbreviated name for Sunday is `abbrev_weekday_names[0]` and the abbreviated name for Saturday is `abbrev_weekday_names[6]`.

**char \*weekday\_names[7]**

The full names of the weekdays; for example, the full name for Sunday is `weekday_names[0]` and the full name for Saturday is `weekday_names[6]`.

**char \*time\_format**

The standard format for times, using the format specifiers supported by `strftime()` and `strptime()` (see `ctime(3V)`).

**char \*sdate\_format**

The standard short format for dates, using the format specifiers supported by `ctime(3V)`.

**char \*dtime\_format**

The standard short format for dates and times together, using the format specifiers supported by `ctime(3V)`.

**char \*am\_string**

The string representing AM.

**char \*pm\_string**

The string representing PM.

**char \*ldate\_format**

The standard long format for dates, using the format specifiers supported by `ctime(3V)`.

The values for the members in the C locale are:

<code>abbrev_month_names</code>	Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec
<code>month_names</code>	January, February, March, April, May, June, July, August, September, October, November, December
<code>abbrev_weekday_names</code>	Sun, Mon, Tue, Wed, Thu, Fri, Sat
<code>weekday_names</code>	Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday
<code>time_format</code>	%H:%M:%S

<b>sdate_format</b>	<b>%m/%d/%y</b>
<b>dtime_format</b>	<b>%a %b %e %T %Z %Y</b>
<b>am_string</b>	<b>AM</b>
<b>pm_string</b>	<b>PM</b>
<b>ldate_format</b>	<b>%A, %B %e, %Y</b>

**FILES**

**/usr/share/lib/locale/LC\_TIME**

standard locale information directory for category LC\_TIME

**SEE ALSO**

**ctime(3V), setlocale(3V)**



## NAME

localeconv – get numeric and monetary formatting conventions

## SYNOPSIS

```
#include <limits.h>
```

```
#include <locale.h>
```

```
struct lconv *localeconv()
```

## DESCRIPTION

`localeconv()` returns a pointer to a structure of type `struct lconv` containing values appropriate for the formatting of numeric quantities (monetary and otherwise) according to the rules of the current locale.

The members of the structure with type `(char *)` are strings; if a string has the value `""`, the value is not available in the current locale or has zero length. The members with type `char` are nonnegative numbers; if any of them have the value `CHAR_MAX` the value is not available in the current locale. The `lconv` structure is defined in `<locale.h>` as follows:

```
struct lconv {
    char    *decimal_point;    /* decimal point character */
    char    *thousands_sep;   /* thousands separator character */
    char    *grouping;        /* grouping of digits */
    char    *int_curr_symbol;  /* international currency symbol */
    char    *currency_symbol; /* local currency symbol */
    char    *mon_decimal_point; /* monetary decimal point character */
    char    *mon_thousands_sep; /* monetary thousands separator */
    char    *mon_grouping;    /* monetary grouping of digits */
    char    *positive_sign;   /* monetary credit symbol */
    char    *negative_sign;   /* monetary debit symbol */
    char    int_frac_digits;  /* intl monetary number of fractional digits */
    char    frac_digits;     /* monetary number of fractional digits */
    char    p_cs_precedes;    /* true if currency symbol precedes credit */
    char    p_sep_by_space;   /* true if space separates c.s. from credit */
    char    n_cs_precedes;    /* true if currency symbol precedes debit */
    char    n_sep_by_space;   /* true if space separates c.s. from debit */
    char    p_sign_posn;      /* position of sign for credit */
    char    n_sign_posn;      /* position of sign for debit */
};
```

The fields of this structure represent:

**decimal\_point**

The decimal-point character used to format non-monetary quantities.

**thousands\_sep**

The character used to separate groups of digits to the left of the decimal-point character in formatted non-monetary quantities.

**grouping**

A string whose elements indicate the size of each group of digits in formatted non-monetary quantities.

**int\_curr\_symbol**

The international currency symbol applicable to the current locale, left-justified within a four-character SPACE-padded field. The character sequences are those specified in: *ISO 4217 Codes for the Representation of Currency and Funds*.

**currency\_symbol**

The local currency symbol applicable to the current locale.

**mon\_decimal\_point**

The decimal-point used to format monetary quantities.

**mon\_thousands\_sep**

The character used to separate groups of digits to the left of the decimal-point character in formatted monetary quantities.

**mon\_grouping**

A string whose elements indicate the size of each group of digits in formatted monetary quantities.

**positive\_sign**

The string used to indicate a nonnegative-valued formatted monetary quantity.

**negative\_sign**

The string used to indicate a negative-valued formatted monetary quantity.

**int\_frac\_digits**

The number of fractional digits (those after the decimal-point) to be displayed in an internationally formatted monetary quantity.

**frac\_digits**

The number of fractional digits (those to the right of the decimal-point) to be displayed in a formatted monetary quantity.

**p\_cs\_precedes**

1 if the **currency\_symbol** precedes the value for a nonnegative formatted monetary quantity; 0 if the **currency\_symbol** succeeds the value for a nonnegative formatted monetary quantity.

**p\_sep\_by\_space**

1 if the **currency\_symbol** is separated by a SPACE from the value for a nonnegative formatted monetary quantity; 0 if the **currency\_symbol** is not separated by a SPACE from the value for a nonnegative formatted monetary quantity.

**n\_cs\_precedes**

1 if the **currency\_symbol** precedes the value for a negative formatted monetary quantity; 0 if the **currency\_symbol** succeeds the value for a negative formatted monetary quantity.

**n\_sep\_by\_space**

1 if the **currency\_symbol** is separated by a SPACE from the value for a negative formatted monetary quantity; 0 if the **currency\_symbol** is not separated by a SPACE from the value for a negative formatted monetary quantity.

**p\_sign\_posn**

A value indicating the positioning of the **positive\_sign** for a nonnegative formatted monetary quantity.

**n\_sign\_posn**

A value indicating the positioning of the **negative\_sign** for a negative formatted monetary quantity.

The elements of **grouping** and **mon\_grouping** are interpreted as follows:

<b>CHAR_MAX</b>	No further grouping is to be performed.
<b>0</b>	The previous element is to be repeatedly used for the remainder of the digits.
<i>other</i>	The value is the number of digits that comprise the current group. The next element is examined to determine the size of the next group of digits to the left of the current group.

The values of **p\_sign\_posn** and **n\_sign\_posn** are interpreted as follows:

<b>0</b>	Parentheses surround the quantity and <b>currency_symbol</b> .
<b>1</b>	The sign string precedes the quantity and <b>currency_symbol</b> .

- 2 The sign string succeeds the quantity and `currency_symbol`.
- 3 The sign string immediately precedes the `currency_symbol`.
- 4 The sign string immediately succeeds the `currency_symbol`.

The values for the members in the C locale are:

<i>field</i>	<i>value</i>
<code>decimal_point</code>	"."
<code>thousands_sep</code>	""
<code>grouping</code>	""
<code>int_curr_symbol</code>	""
<code>currency_symbol</code>	""
<code>mon_decimal_point</code>	""
<code>mon_thousands_sep</code>	""
<code>mon_grouping</code>	""
<code>positive_sign</code>	""
<code>negative_sign</code>	""
<code>int_frac_digits</code>	CHAR_MAX
<code>frac_digits</code>	CHAR_MAX
<code>p_cs_precedes</code>	CHAR_MAX
<code>p_sep_by_space</code>	CHAR_MAX
<code>n_cs_precedes</code>	CHAR_MAX
<code>n_sep_by_space</code>	CHAR_MAX
<code>p_sign_posn</code>	CHAR_MAX
<code>n_sign_posn</code>	CHAR_MAX

#### RETURN VALUES

`localeconv()` returns a pointer to `struct lconv` (see NOTES).

#### FILES

`/usr/share/lib/locale/LC_MONETARY`

standard locale information directory for category `LC_MONETARY`

`/usr/share/lib/locale/LC_NUMERIC`

standard locale information directory for category `LC_NUMERIC`

#### SEE ALSO

`printf(3V)`, `scanf(3V)`, `setlocale(3V)`

#### NOTES

`localeconv()` does not modify the `struct lconv` to which it returns a pointer, but subsequent calls to `setlocale(3V)` with categories `LC_ALL`, `LC_MONETARY`, or `LC_NUMERIC` may overwrite the contents of the structure.

## NAME

lockf – record locking on files

## SYNOPSIS

```
#include <unistd.h>

int lockf(fd, cmd, size)
int fd, cmd;
long size;
```

## DESCRIPTION

**lockf()** places, removes, and tests for exclusive locks on sections of files. These locks are either advisory or mandatory depending on the mode bits of the file. The lock is mandatory if the set-GID bit (**S\_ISGID**) is set and the group execute bit (**S\_IXGRP**) is clear (see **stat(2V)** for information about mode bits). Otherwise, the lock is advisory.

If a process holds a mandatory exclusive lock on a segment of a file, both read and write operations block until the lock is removed (see **WARNINGS**).

An advisory lock does not affect read and write access to the locked segment. Advisory locks may be used by cooperating processes checking for locks using **F\_GETLCK** and voluntarily observing the indicated read and write restrictions.

A locking call on an already locked file section fails, returning an error value or putting the call to sleep until that file section is unlocked. All the locks on a process are removed when that process terminates. See **fcntl(2V)** for more information about record locking.

*fd* is an open file descriptor. It must have **O\_WRONLY** or **O\_RDWR** permission for a successful locking call.

*cmd* is a control value which specifies the action to be taken. The accepted values for *cmd* are defined in **<unistd.h>** as follows:

```
#define F_ULOCK    0    /* Unlock a previously locked section */
#define F_LOCK     1    /* Lock a section for exclusive use */
#define F_TLOCK    2    /* Test and lock a section (non-blocking) */
#define F_TEST     3    /* Test section for other process' locks */
```

**F\_TEST** returns **-1** and sets **errno** to **EACCES** if a lock by another process already exists on the specified section. Otherwise, it returns **0**. **F\_LOCK** and **F\_TLOCK** lock available file sections. **F\_ULOCK** removes locks from file sections.

All other values of *cmd* are reserved for future applications and, until implemented, return an error.

*size* is the number of contiguous bytes to be locked or unlocked. The resource to be locked starts at the current offset in the file and extends forward *size* bytes if *size* is positive, and extends backward *size* bytes (the preceding bytes up to but not including the current offset) if *size* is negative. If *size* is zero, the section from the current offset through the largest file offset is locked (that is, from the current offset through the present or any future EOF). An area need not be allocated to the file to be locked, such a lock may exist after the EOF.

Sections locked with **F\_LOCK** or **F\_TLOCK** may contain all or part of an already locked section. They may also be partially or completely contained by an already locked section. Where these overlapping or adjacent locked sections occur, they are combined into a single section. If the table of active locks is full, a lock request requiring an additional table entry fails and an error value is returned.

**F\_LOCK** and **F\_TLOCK** differ only in their response to requests for unavailable resources. If a section is already locked, **F\_LOCK** directs the calling process to sleep until the resource is available, **F\_TLOCK** directs the function to return **-1** and set **errno** to **EACCES** (see **ERRORS**).

When a `F_ULOCK` request releases part of a section with overlapping locks, the remaining section or sections retain the lock. If `F_ULOCK` removes the center of a locked section, the two separate locked sections remain, but an additional element is required in the table of active locks. If this table is full, `errno` is set to `ENOLCK` and the requested section is not released.

The danger of a deadlock exists when a process controlling a locked resource is put to sleep by requesting an unavailable resource. To avoid this danger, `lockf()` and `fcntl()` scan for this conflict before putting a locked resource to sleep. If a deadlock would result, an error value is returned.

The sleep process can be interrupted with any signal. `alarm(3V)` may be used to provide a timeout facility where needed.

#### RETURN VALUES

`lockf()` returns:

- 0        on success.
- 1       on failure and sets `errno` to indicate the error.

#### ERRORS

- `EACCES`        *cmd* is `F_TLOCK` or `F_TEST` and the section is already locked by another process.  
Note: In future, `lockf()` may generate `EAGAIN` under these conditions, so applications testing for `EACCES` should also test for `EAGAIN`.
- `EBADF`        *fd* is not a valid open descriptor.  
*cmd* is `F_LOCK` or `F_TLOCK` and the process does not have write permission on the file.
- `EDEADLK`      *cmd* is `F_LOCK` and a deadlock would occur.
- `EINTR`        *cmd* is `F_LOCK` and a signal interrupted the process while it was waiting to complete the lock.
- `ENOLCK`      *cmd* is `F_LOCK`, `F_TLOCK`, or `F_ULOCK` and there are no more file lock entries available.

#### SEE ALSO

`chmod(2V)`, `fcntl(2V)`, `flock(2)`, `fork(2V)`, `alarm(3V)`, `lockd(8C)`

#### WARNINGS

Mandatory record locks are dangerous. If a runaway or otherwise out-of-control process should hold a mandatory lock on a file critical to the system and fail to release that lock, the entire system could hang or crash. For this reason, mandatory record locks may be removed in a future SunOS release. Use advisory record locking whenever possible.

#### NOTES

A child process does not inherit locks from its parent on `fork(2V)`.

#### BUGS

`lockf()` locks do not interact in any way with locks granted by `flock()`, but are compatible with locks granted by `fcntl()`.

## NAME

`lsearch`, `lfind` – linear search and update

## SYNOPSIS

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

char *lsearch (key, base, nelp, width, compar)
char *key;
char *base;
unsigned int *nelp;
unsigned int width;
int (*compar)();

char *lfind (key, base, nelp, width, compar)
char *key;
char *base;
unsigned int *nelp;
unsigned int width;
int (*compar)();
```

## DESCRIPTION

`lsearch()` is a linear search routine generalized from Knuth (6.1) Algorithm S. It returns a pointer into a table indicating where a datum may be found. If the datum does not occur, it is added at the end of the table. *key* points to the datum to be sought in the table. *base* points to the first element in the table. *nelp* points to an integer containing the current number of elements in the table. The integer is incremented if the datum is added to the table. *compar* is the name of the comparison function which the user must supply (`strcmp()`, for example). It is called with two arguments that point to the elements being compared. The function must return zero if the elements are equal and non-zero otherwise.

`lfind()` is the same as `lsearch()` except that if the datum is not found, it is not added to the table. Instead, a NULL pointer is returned.

## NOTES

The pointers to the key and the element at the base of the table should be of type pointer-to-element, and cast to type pointer-to-character.

The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

Although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

## EXAMPLE

This fragment will read in  $\leq$  TABSIZE strings of length  $\leq$  ELSIZE and store them in a table, eliminating duplicates.

```
#include <stdio.h>
#include <search.h>
#define
TABSIZE 50
#define
ELSIZE 120
char line[ELSIZE], tab[TABSIZE][ELSIZE], *lsearch( );
unsigned nel = 0;
int strcmp( );
...
while (fgets(line,
ELSIZE, stdin) != NULL &&
```

**nel < TABSIZE)**

**(void) lsearch(line, (char \*)tab, &nel, ELSIZE, strcmp);**

...

**SEE ALSO**

**bsearch(3), hsearch(3), tsearch(3)**

**DIAGNOSTICS**

If the searched for datum is found, both **lsearch()** and **lfind()** return a pointer to it. Otherwise, **lfind()** returns NULL and **lsearch()** returns a pointer to the newly added element.

**BUGS**

Undefined results can occur if there is not enough room in the table to add a new item.

**NAME**

`madvise` – provide advice to VM system

**SYNOPSIS**

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

int madvise(addr, len, advice)
caddr_t addr;
size_t len;
int advice;
```

**DESCRIPTION**

`madvise()` advises the kernel that a region of user mapped memory in the range `[addr, addr + len)` will be accessed following a type of pattern. The kernel uses this information to optimize the procedure for manipulating and maintaining the resources associated with the specified mapping range.

Values for *advice* are defined in `<sys/mman.h>` as:

```
#define MADV_NORMAL    0x0    /* No further special treatment */
#define MADV_RANDOM    0x1    /* Expect random page references */
#define MADV_SEQUENTIAL 0x2 /* Expect sequential page references */
#define MADV_WILLNEED  0x3    /* Will need these pages */
#define MADV_DONTNEED  0x4    /* Don't need these pages */
```

**MADV\_NORMAL**

The default system characteristic where accessing memory within the address range causes the system to read data from the mapped file. The kernel reads all data from files into pages which are retained for a period of time as a “cache”. System pages can be a scarce resource, so the kernel steals pages from other mappings when needed. This is a likely occurrence but only adversely affects system performance if a large amount of memory is accessed.

**MADV\_RANDOM**

Tells the kernel to read in a minimum amount of data from a mapped file when doing any single particular access. Normally when an address of a mapped file is accessed, the system tries to read in as much data from the file as reasonable, in anticipation of other accesses within a certain locality.

**MADV\_SEQUENTIAL**

Tells the system that addresses in this range are likely to only be accessed once, so the system will free the resources used to map the address range as quickly as possible. This is used in the `cat(1V)` and `cp(1)` utilities.

**MADV\_WILLNEED**

Tells the system that a certain address range is definitely needed, so the kernel will read the specified range into memory immediately. This might be beneficial to programs who want to minimize the time it takes to access memory the first time since the kernel would need to read in from the file.

**MADV\_DONTNEED**

Tells the kernel that the specified address range is no longer needed, so the system immediately frees the resources associated with the address range.

`madvise()` should be used by programs that have specific knowledge of their access patterns over a memory object (for example, a mapped file) and wish to increase system performance.

**RETURN VALUES**

`madvise()` returns:

- 0        on success.
- 1       on failure and sets `errno` to indicate the error.



**ERRORS****EINVAL**

*addr* is not a multiple of the page size as returned by `getpagesize(2)`.  
The length of the specified address range is less than or equal to 0.

*advice* was invalid.

**EIO**

An I/O error occurred while reading from or writing to the file system.

**ENOMEM**

Addresses in the range [*addr*, *addr + len*) are outside the valid range for the address space of a process, or specify one or more pages that are not mapped.

**SEE ALSO****mctl(2), mmap(2)**

## NAME

malloc, free, realloc, calloc, cfree, memalign, valloc, mallocmap, mallopt, mallinfo, malloc\_debug, malloc\_verify, alloca – memory allocator

## SYNOPSIS

```
#include <malloc.h>

char *malloc(size)
unsigned size;

int free(ptr)
char *ptr;

char *realloc(ptr, size)
char *ptr;
unsigned size;

char *calloc(nelem, elsize)
unsigned nelem, elsize;

int cfree(ptr)
char *ptr;

char *memalign(alignment, size)
unsigned alignment;
unsigned size;

char *valloc(size)
unsigned size;

void mallocmap()

int mallopt(cmd, value)
int cmd, value;

struct mallinfo mallinfo()

#include <alloca.h>

char *alloca(size)
int size;
```

## SYSTEM V SYNOPSIS

```
#include <malloc.h>

void *malloc(size)
size_t size;

void free(ptr)
void *ptr;

void *realloc(ptr, size)
void *ptr;
size_t size;

void *calloc(nelem, elsize)
size_t nelem;
size_t elsize;

void *memalign(alignment, size)
size_t alignment;
size_t size;

void *valloc(size)
size_t size;
```

The XPG2 versions of the functions listed in this section are declared as they are in SYNOPSIS above, except `free()`, which is declared as:

```
void free(ptr)
char *ptr;
```

#### DESCRIPTION

These routines provide a general-purpose memory allocation package. They maintain a table of free blocks for efficient allocation and coalescing of free storage. When there is no suitable space already free, the allocation routines call `sbrk()` (see `brk(2)`) to get more memory from the system.

Each of the allocation routines returns a pointer to space suitably aligned for storage of any type of object. Each returns a NULL pointer if the request cannot be completed (see **DIAGNOSTICS**).

`malloc()` returns a pointer to a block of at least *size* bytes, which is appropriately aligned.

`free()` releases a previously allocated block. Its argument is a pointer to a block previously allocated by `malloc()`, `calloc()`, `realloc()`, `valloc()`, or `memalign()`.

`realloc()` changes the size of the block referenced by *ptr* to *size* bytes and returns a pointer to the (possibly moved) block. The contents will be unchanged up to the lesser of the new and old sizes. If unable to honor a reallocation request, `realloc()` leaves its first argument unaltered. For backwards compatibility, `realloc()` accepts a pointer to a block freed since the most recent call to `malloc()`, `calloc()`, `realloc()`, `valloc()`, or `memalign()`. Note: using `realloc()` with a block freed *before* the most recent call to `malloc()`, `calloc()`, `realloc()`, `valloc()`, or `memalign()` is an error.

`calloc()` uses `malloc()` to allocate space for an array of *nelem* elements of size *elsize*, initializes the space to zeros, and returns a pointer to the initialized block. The block can be freed with `free()` or `cfree()`.

`memalign()` allocates *size* bytes on a specified alignment boundary, and returns a pointer to the allocated block. The value of the returned address is guaranteed to be an even multiple of *alignment*. Note: the value of *alignment* must be a power of two, and must be greater than or equal to the size of a word.

`valloc(size)` is equivalent to `memalign(getpagesize(), size)`.

`mallocmap()` prints a map of the heap to the standard output. `mallocmap()` prints each block's address, size (in bytes) and status (free or busy). A block must have a size that is no larger than the current extent of the heap.

`mallopt()` allows quick allocation of small blocks of memory. `mallopt()` tells subsequent calls to `malloc()` to allocate *holding blocks* containing small blocks. Under this small block algorithm, a request to `malloc()` for a small block of memory returns a pointer to one of the pre-allocated small blocks. Different holding blocks are created as needed for different sizes of small blocks.

*cmd* may be one of the following values, defined in `<malloc.h>`:

- |                 |  |
|-----------------|--|
| <b>M_MXFAST</b> | Set the maximum size of blocks to be allocated using the small block algorithm ( <i>maxfast</i> ) to <i>value</i> . The algorithm allocates all blocks smaller than <i>maxfast</i> in large groups and then does them out very quickly. Initially, <i>maxfast</i> is 0 and the small block algorithm is disabled.  |
| <b>M_NLBLKS</b> | Set the number of small blocks in a holding block ( <i>numlblks</i> ) to <i>value</i> . The holding blocks each contain <i>numlblks</i> blocks. <i>numlblks</i> must be greater than 1. The default value for <i>numlblks</i> is 100.  |
| <b>M_GRAIN</b>  | Set the granularity for small block requests ( <i>grain</i> ) to <i>value</i> . The sizes of all blocks smaller than <i>maxfast</i> are rounded up to the nearest multiple of <i>grain</i> . <i>grain</i> must be greater than 0. The default value of <i>grain</i> is the smallest number of bytes which will allow alignment of any data type. When <i>grain</i> is set, <i>value</i> is rounded up to a multiple of this default. |

**M\_KEEP** Preserve data in a freed block until the next `malloc()`, `realloc()`, or `calloc()`. This option is provided only for compatibility with the old version of `malloc()` and is not recommended.

`mallocpt()` may be called repeatedly, but may not be called after the first small block is allocated.

`mallinfo()` can be used during program development to determine the best settings for the parameters set by `mallocpt()`. Do not call `mallinfo()` until after a call to `malloc()`. `mallinfo()` provides information describing space usage. It returns a `mallinfo` structure, defined in `<malloc.h>` as:

```
struct mallinfo {
    int arena;      /* total space in arena */
    int ordblks;   /* number of ordinary blocks */
    int smlblks;   /* number of small blocks */
    int hblks;     /* number of holding blocks */
    int hblkhd;    /* space in holding block headers */
    int usmlblks;  /* space in small blocks in use */
    int fsmblks;   /* space in free small blocks */
    int uordblks;  /* space in ordinary blocks in use */
    int fordblks;  /* space in free ordinary blocks */
    int keepcost;  /* cost of enabling keep option */

    int mxfast;    /* max size of small blocks */
    int nblks;     /* number of small blocks in a holding block */
    int grain;     /* small block rounding factor */
    int uordbytes; /* space (including overhead) allocated in ord. blks */
    int allocated; /* number of ordinary blocks allocated */
    int treeoverhead; /* bytes used in maintaining the free tree */
};
```

`alloca()` allocates *size* bytes of space in the stack frame of the caller, and returns a pointer to the allocated block. This temporary space is automatically freed when the caller returns. Note that if the allocated block is beyond the current stack limit, the resulting behavior is undefined.

`malloc()`, `realloc()`, `memalign()` and `valloc()` return a non-NULL pointer if *size* is 0, and `calloc()` returns a non-NULL pointer if *nelem* or *elsize* is 0, but these pointers should *not* be dereferenced.

Note: Always cast the value returned by `malloc()`, `realloc()`, `calloc()`, `memalign()`, `valloc()` or `alloca()`.

#### SYSTEM V DESCRIPTION

The XPG2 versions of `malloc()`, `realloc()`, `memalign()` and `valloc()` return NULL if *size* is 0. The XPG2 version of `calloc()` returns NULL if *nelem* or *elsize* is 0.

#### RETURN VALUES

On success, `malloc()`, `calloc()`, `realloc()`, `memalign()`, `valloc()` and `alloca()` return a pointer to space suitably aligned for storage of any type of object. On failure, they return NULL.

`free()` and `cfree()` return:

1 on success.

0 on failure and set `errno` to indicate the error.

`mallocpt()` returns 0 on success. If `mallocpt()` is called after the allocation of a small block, or if *cmd* or *value* is invalid, it returns a non-zero value.

`mallinfo()` returns a `struct mallinfo`.

**SYSTEM V RETURN VALUES**

If *size* is 0, the XPG2 versions of **malloc()**, **realloc()**, **memalign()** and **valloc()** return NULL.

If *nelem* or *elsize* is 0, the XPG2 version of **calloc()** returns NULL.

**free()** does not return a value.

**ERRORS**

**malloc()**, **calloc()**, **realloc()**, **valloc()**, **memalign()**, **cfree()**, and **free()** will each fail if one or more of the following are true:

**EINVAL** An invalid argument was specified.

The value of *ptr* passed to **free()**, **cfree()**, or **realloc()** was not a pointer to a block previously allocated by **malloc()**, **calloc()**, **realloc()**, **valloc()**, or **memalign()**.

The allocation heap is found to have been corrupted. More detailed information may be obtained by enabling range checks using **malloc\_debug()**.

**ENOMEM** *size* bytes of memory could not be allocated.

**FILES**

**/usr/lib/debug/malloc.o** diagnostic versions of **malloc()** routines.

**/usr/lib/debug/mallocmap.o** routines to print a map of the heap.

**SEE ALSO**

**csh(1)**, **ld(1)**, **brk(2)**, **getrlimit(2)**, **sigvec(2)**, **sigstack(2)**

Stephenson, C.J., *Fast Fits*, in *Proceedings of the ACM 9th Symposium on Operating Systems*, SIGOPS *Operating Systems Review*, vol. 17, no. 5, October 1983.

*Core Wars*, in *Scientific American*, May 1984.

**DIAGNOSTICS**

More detailed diagnostics can be made available to programs using **malloc()**, **calloc()**, **realloc()**, **valloc()**, **memalign()**, **cfree()**, and **free()**, by including a special relocatable object file at link time (see FILES). This file also provides routines for control of error handling and diagnosis, as defined below. Note: these routines are *not* defined in the standard library.

```
int malloc_debug(level)
```

```
int level;
```

```
int malloc_verify()
```

**malloc\_debug()** sets the level of error diagnosis and reporting during subsequent calls to **malloc()**, **calloc()**, **realloc()**, **valloc()**, **memalign()**, **cfree()**, and **free()**. The value of *level* is interpreted as follows:

Level 0 **malloc()**, **calloc()**, **realloc()**, **valloc()**, **memalign()**, **cfree()**, and **free()** behave the same as in the standard library.

Level 1 The routines abort with a message to the standard error if errors are detected in arguments or in the heap. If a bad block is encountered, its address and size are included in the message.

Level 2 Same as level 1, except that the entire heap is examined on every call to the above routines.

**malloc\_debug()** returns the previous error diagnostic level. The default level is 1.

**malloc\_verify()** attempts to determine if the heap has been corrupted. It scans all blocks in the heap (both free and allocated) looking for strange addresses or absurd sizes, and also checks for inconsistencies in the free space table. **malloc\_verify()** returns 1 if all checks pass without error, and otherwise returns 0. The checks can take a significant amount of time, so it should not be used indiscriminately.

**WARNINGS**

**alloca()** is machine-, compiler-, and most of all, system-dependent. Its use is strongly discouraged. See **getrlimit(2)**, **sigvec(2)**, **sigstack(2)**, **csh(1)**, and **ld(1)**.

**NOTES**

Because **malloc()**, **realloc()**, **memalign()** and **valloc()** return a non-NULL pointer if *size* is 0, and **calloc()** returns a non-NULL pointer if *nelem* or *elsize* is 0, a zero size need not be treated as a special case if it should be passed to these functions unpredictably. Also, the pointer returned by these functions may be passed to subsequent invocations of **realloc()**.

**SYSTEM V NOTES**

The XPG2 versions of the allocation routines return NULL when passed a zero size (see **SYSTEM V DESCRIPTION** above).

**BUGS**

Since **realloc()** accepts a pointer to a block freed since the last call to **malloc()**, **calloc()**, **realloc()**, **valloc()**, or **memalign()**, a degradation of performance results. The semantics of **free()** should be changed so that the contents of a previously freed block are undefined.

**NAME**

`mblen`, `mbstowcs`, `mbtowc`, `wcstombs`, `wctomb` – multibyte character handling

**SYNOPSIS**

```
#include <stdlib.h>

int mblen(s, n)
char *s;
size_t n;

size_t mbstowcs(s, pwcs, n)
char *s;
wchar_t *pwcs;
size_t n;

int mbtowc(pwc, s, n)
wchar_t *pwc;
char *s;
size_t n;

int wcstombs(s, pwcs, n)
char *s;
wchar_t *pwcs;
size_t n;

int wctomb(s, wchar)
char *s;
wchar_t wcar;
```

**DESCRIPTION**

The behavior of these functions is affected by the `LC_CTYPE` category of the program's locale. For a stat-dependent encoding, each function is placed into its initial state by a call for which its character pointer argument, *s*, is a NULL pointer. Subsequent calls with *s* as other than a NULL pointer cause the internal state of the function to be altered as necessary. A call with a *s* as a NULL pointer causes these functions to return a nonzero value if encodings have state dependency, and zero otherwise. After the `LC_CTYPE` category is changed, the shift state of these functions is indeterminate.

If *s* is not a NULL pointer, these functions work as follows:

**mblen()**

Determines the number of bytes comprising the multibyte character pointed to by *s*.

**mbstowcs()**

Converts a sequence of multibyte characters that begins in the initial shift state from the array pointed to by *s* into a sequence of corresponding codes and stores no more than *n* codes into the array pointed to by *pwcs*. No multibyte characters that follow a null character (which is converted into a code with value zero) will be examined or converted. Each multibyte character is converted as if by a call to `mbtowc()`, except that the shift state of `mbtowc()` is not affected.

No more than *n* elements will be modified in the array pointed to by *pwcs*. If copying takes place between objects that overlap, the behavior is undefined.

**mbtowc()**

Determines the number of bytes that comprise the multibyte character pointed to by *s*. `mbtowc()` then determines the code for value of type `wchar_t` that corresponds to that multibyte character. The value of the code corresponding to the null character is zero. If the multibyte character is valid and *pwc* is not a null pointer, `mbtowc()` stores the code in the object pointed to by *pwc*. At most *n* bytes of the array pointed to by *s* will be examined.

**wcstowcs()**

Converts a sequence of codes that correspond to multibyte characters from the array pointed to by *pwcs* into a sequence of multibyte characters that begins in the initial shift state and stores these multibyte characters into the array pointed to by *s*, stopping if a multibyte character would exceed the limit of *n* total bytes or if a null character is stored. Each code is converted as if by a call to **wctomb()**, except that the shift state of **wctomb()** is not affected.

**wctomb()**

Determines the number of bytes needed to represent the multibyte character corresponding to the code whose value is *wchar* (including any change in shift state). **wctomb()** stores the multibyte character representation in the array object pointed to by *s* (if *s* is not a null pointer). At most, **MB\_CUR\_MAX** characters are stored. If the value of *wchar* is zero, **wctomb()** is left in the initial shift state.

**RETURN VALUES**

If *s* is a null pointer, **mblen()**, **mbtowc()**, and **wctomb()** return a nonzero or zero value, if multibyte character encodings, respectively, do or do not have state dependent encodings.

If *s* is not a null pointer, **mblen()** and **mbtowc()** either return 0 (if *s* points to the null character), or return the number of bytes that comprise the converted multibyte character (if the next *n* or fewer bytes form a valid multibyte character), or return -1 (if they do not form a valid multibyte character).

In no case will the value returned by **mbtowc()** be greater than *n* or the value of the **MB\_CUR\_MAX** macro. If *s* is not a null pointer, **wctomb()** returns -1 (if the value does not correspond to a valid multibyte character), or returns the number of bytes that comprise the multibyte character corresponding to *wchar*.

If an invalid multibyte character is encountered, **mbstowcs()** and **wcstombs()** return **(size\_t) -1**. Otherwise, they return the number of bytes modified, not including a terminating null character, if any.



**NAME**

memory, memccpy, memchr, memcmp, memcpy, memset – memory operations

**SYNOPSIS**

```
#include <memory.h>

char *memccpy(s1, s2, c, n)
char *s1, *s2;
int c, n;

char *memchr(s, c, n)
char *s;
int c, n;

int memcmp(s1, s2, n)
char *s1, *s2;
int n;

char *memcpy(s1, s2, n)
char *s1, *s2;
int n;

char *memset(s, c, n)
char *s;
int c, n;
```

**DESCRIPTION**

These functions operate as efficiently as possible on memory areas (arrays of characters bounded by a count, not terminated by a null character). They do not check for the overflow of any receiving memory area.

**memccpy()** copies characters from memory area *s2* into *s1*, stopping after the first occurrence of character *c* has been copied, or after *n* characters have been copied, whichever comes first. It returns a pointer to the character after the copy of *c* in *s1*, or a NULL pointer if *c* was not found in the first *n* characters of *s2*.

**memchr()** returns a pointer to the first occurrence of character *c* in the first *n* characters of memory area *s*, or a NULL pointer if *c* does not occur.

**memcmp()** compares its arguments, looking at the first *n* characters only, and returns an integer less than, equal to, or greater than 0, according as *s1* is lexicographically less than, equal to, or greater than *s2*.

**memcpy()** copies *n* characters from memory area *s2* to *s1*. It returns *s1*.

**memset()** sets the first *n* characters in memory area *s* to the value of character *c*. It returns *s*.

**NOTES**

For user convenience, all these functions are declared in the **<memory.h>** header file.

**BUGS**

**memcmp()** uses native character comparison, which is signed on some machines and unsigned on other machines. Thus the sign of the value returned when one of the characters has its high-order bit set is implementation-dependent.

Character movement is performed differently in different implementations. Thus overlapping moves may yield surprises.

**NAME**

**mktemp**, **mkstemp** – make a unique file name

**SYNOPSIS**

```
char *mktemp(template)
```

```
char *template;
```

```
mkstemp(template)
```

```
char *template;
```

**DESCRIPTION**

**mktemp()** creates a unique file name, typically in a temporary filesystem, by replacing *template* with a unique file name, and returns the address of *template*. The string in *template* should contain a file name with six trailing Xs; **mktemp()** replaces the Xs with a letter and the current process ID. The letter will be chosen so that the resulting name does not duplicate an existing file. **mkstemp()** makes the same replacement to the template but returns a file descriptor for the template file open for reading and writing. **mkstemp()** avoids the race between testing whether the file exists and opening it for use.

**Notes:**

- **mktemp()** and **mkstemp()** actually *change* the template string which you pass; this means that you cannot use the same template string more than once — you need a fresh template for every unique file you want to open.
- When **mktemp()** or **mkstemp()** are creating a new unique filename they check for the prior existence of a file with that name. This means that if you are creating more than one unique filename, it is bad practice to use the same root template for multiple invocations of **mktemp()** or **mkstemp()**.

**SEE ALSO**

**getpid(2V)**, **open(2V)**, **tmpfile(3S)**, **tmpnam(3S)**

**DIAGNOSTICS**

**mkstemp()** returns an open file descriptor upon success. It returns **-1** if no suitable file could be created.

**mktemp()** assigns the null string to *template* when it cannot create a unique name.

**BUGS**

It is possible to run out of letters.

**NAME**

**mlock, munlock** – lock (or unlock) pages in memory

**SYNOPSIS**

```
#include <sys/types.h>
int mlock(addr, len) caddr_t addr; size_t len;

int munlock(addr, len)
caddr_t addr;
size_t len;
```

**DESCRIPTION**

**mlock()** uses the mappings established for the address range [*addr*, *addr + len*) to identify memory object pages to be locked in memory. If the page identified by a mapping changes, such as occurs when a copy of a writable **MAP\_PRIVATE** page is made upon the first store, the lock will be transferred to the newly copied private page.

**munlock()** removes locks established with **mlock()**.

A given page may be locked multiple times by executing an **mlock()** through different mappings. That is, if two different processes lock the same page then the page will remain locked until both processes remove their locks. However, within a given mapping, page locks do not nest – multiple **mlock()** operations on the same address in the same process will all be removed with a single **munlock()**. Of course, a page locked in one process and mapped in another (or visible through a different mapping in the locking process) is still locked in memory. This fact can be used to create applications that do nothing other than lock important data in memory, thereby avoiding page I/O faults on references from other processes in the system.

If the mapping through which an **mlock()** has been performed is removed, an **munlock()** is implicitly performed. An **munlock()** is also performed implicitly when a page is deleted through file removal or truncation.

Locks established with **mlock()** are not inherited by a child process after a **fork(2V)**.

Due to the impact on system resources, the use of **mlock()** and **munlock()** is restricted to the super-user. Attempts to **mlock()** more memory than a system-specific limit will fail.

**RETURN VALUES**

**mlock()** and **munlock()** return:

- 0        on success.
- 1       on failure and set **errno** to indicate the error.

**ERRORS**

- EAGAIN**        (**mlock()** only.) Some or all of the memory identified by the range [*addr*, *addr + len*) could not be locked due to insufficient system resources.
- EINVAL**        *addr* is not a multiple of the page size as returned by **getpagesize(2)**.
- ENOMEM**        Addresses in the range [*addr*, *addr + len*) are invalid for the address space of a process, or specify one or more pages which are not mapped.
- EPERM**         The process's effective user ID is not super-user.

**SEE ALSO**

**fork(2V)**, **mctl(2)**, **mlockall(3)**, **mmap(2)**, **munmap(2)**

**NAME**

**mlockall**, **munlockall** – lock (or unlock) address space

**SYNOPSIS**

```
#include <sys/mman.h>
```

```
int mlockall(flags)
```

```
int flags;
```

```
int munlockall()
```

**DESCRIPTION**

**mlockall()** locks all pages mapped by an address space in memory. The value of *flags* determines whether the pages to be locked are simply those currently mapped by the address space, those that will be mapped in the future, or both. *flags* is built from the options defined in `<sys/mman.h>` as:

```
#define MCL_CURRENT    0x1    /* lock current mappings */
```

```
#define MCL_FUTURE    0x2    /* lock future mappings */
```

If **MCL\_FUTURE** is specified to **mlockall()**, then as mappings are added to the address space (or existing mappings are replaced) they will also be locked, provided sufficient memory is available.

Mappings locked via **mlockall()** with any option may be explicitly unlocked with a **munlock()** call.

**munlockall()** removes address space locks and locks on mappings in the address space.

All conditions and constraints on the use of locked memory as exist for **mlock()** apply to **mlockall()**.

**RETURN VALUES**

**mlockall()** and **munlockall()** return:

0 on success.

-1 on failure and set **errno** to indicate the error.

**ERRORS**

**EAGAIN** (**mlockall()** only.) Some or all of the memory in the address space could not be locked due to sufficient resources.

**EINVAL** *flags* contains values other than **MCL\_CURRENT** and **MCL\_FUTURE**.

**EPERM** The process's effective user ID is not super-user.

**SEE ALSO**

**mctl(2)**, **mlock(3)**, **mmap(2)**

**NAME**

monitor, monstartup, moncontrol – prepare execution profile

**SYNOPSIS**

```
#include <a.out.h>

monitor(lowpc, highpc, buffer, bufsize, nfunc)
int (*lowpc)(), (*highpc)();
short buffer[];

monstartup(lowpc, highpc)
int (*lowpc)(), (*highpc)();

moncontrol(mode)
```

**DESCRIPTION**

There are two different forms of monitoring available. An executable program created by 'cc -p' automatically includes calls for the **prof(1)** monitor, and includes an initial call with default parameters to its start-up routine **monstartup**. In this case, **monitor()** need not be called explicitly, except to gain fine control over **profil(2)** buffer allocation. An executable program created by 'cc -pg' automatically includes calls for the **gprof(1)** monitor.

**monstartup()** is a high-level interface to **profil(2)**. *lowpc* and *highpc* specify the address range that is to be sampled; the lowest address sampled is that of *lowpc* and the highest is just below *highpc*. **monstartup()** allocates space using **sbrk** (see **brk(2)**) and passes it to **monitor()** (as described below) to record a histogram of program-counter values, and calls to certain functions. Only calls to functions compiled with 'cc -p' are recorded.

On Sun-2, Sun-3, and Sun-4 systems, an entire program can be profiled with:

```
extern etext();
...
monstartup(N_TXTOFF(0), etext);
```

On Sun386i systems, the equivalent code sequence is:

```
extern etext();
extern _start();
...
monstartup(_start, etext);
```

**etext** lies just above all the program text, see **end(3)**.

To stop execution monitoring and post results to the file **mon.out**, use:

```
monitor(0);
```

**prof(1)** can then be used to examine the results.

**moncontrol()** is used to selectively control profiling within a program. This works with both **prof(1)** and **gprof(1)**. Profiling begins when the program starts. To stop the collection of profiling statistics, use:

```
moncontrol(0)
```

To resume the collection of statistics, use:

```
moncontrol(1)
```

This allows you to measure the cost of particular functions. Note: an output file is be produced upon program exit, regardless of the state of **moncontrol**.

**monitor()** is a low level interface to **profil(2)**. *lowpc* and *highpc* are the addresses of two functions; *buffer* is the address of a (user supplied) array of *bufsize* short integers. At most *nfunc* call counts can be kept.

For the results to be significant, especially where there are small, heavily used routines, it is suggested that the buffer be no more than a few times smaller than the range of locations sampled. **monitor()** divides the buffer into space to record the histogram of program counter samples over the range *lowpc* to *highpc*, and space to record call counts of functions compiled with the **cc -p**.

To profile the entire program on Sun-2, Sun-3, and Sun-4 systems using the low-level interface to **profil(2)**, it is sufficient to use

```
extern etext();
...
monitor(N_TXTOFF(0), etext, buf, bufsize, nfunc);
```

On Sun386i systems, the equivalent calls are:

```
extern etext();
extern _start();
...
monitor(_start, etext, buf, bufsize, nfunc);
```

**FILES**

**mon.out**

**SEE ALSO**

**cc(1V)**, **prof(1)**, **gprof(1)**, **brk(2)**, **profil(2)**, **end(3)**

**NAME**

mp, madd, msub, mult, mdiv, mcmp, min, mout, pow, gcd, rpow, itom, xtom, mtox, mfree – multiple precision integer arithmetic

**SYNOPSIS**

```
#include <mp.h>

madd(a, b, c)
MINT *a, *b, *c;

msub(a, b, c)
MINT *a, *b, *c;

mult(a, b, c)
MINT *a, *b, *c;

mdiv(a, b, q, r)
MINT *a, *b, *q, *r;

mcmp(a,b)
MINT *a, *b;

min(a)
MINT *a;

mout(a)
MINT *a;

pow(a, b, c, d)
MINT *a, *b, *c, *d;

gcd(a, b, c)
MINT *a, *b, *c;

rpow(a, n, b)
MINT *a, *b;
short n;

msqrt(a, b, r)
MINT *a, *b, *r;

sdiv(a, n, q, r)
MINT *a, *q;
short n, *r;

MINT *itom(n)
short n;

MINT *xtom(s)
char *s;

char *mtox(a)
MINT *a;

void mfree(a)
MINT *a;
```

**DESCRIPTION**

These routines perform arithmetic on integers of arbitrary length. The integers are stored using the defined type MINT. Pointers to a MINT should be initialized using the function `itom()`, which sets the initial value to `n`. Alternatively, `xtom()` may be used to initialize a MINT from a string of hexadecimal digits. `mfree()` may be used to release the storage allocated by the `itom()` and `xtom()` routines.

**madd()**, **msub()** and **mult()** assign to their third arguments the sum, difference, and product, respectively, of their first two arguments. **mdiv()** assigns the quotient and remainder, respectively, to its third and fourth arguments. **sdiv()** is like **mdiv()** except that the divisor is an ordinary integer. **msqrt** produces the square root and remainder of its first argument. **ncmp()** compares the values of its arguments and returns 0 if the two values are equal, a value greater than 0 if the first argument is greater than the second, and a value less than 0 if the second argument is greater than the first. **rpow** raises  $a$  to the  $n$ th power and assigns this value to  $b$ . **pow()** raises  $a$  to the  $b$ th power, reduces the result modulo  $c$  and assigns this value to  $d$ . **min()** and **mout()** do decimal input and output. **gcd()** finds the greatest common divisor of the first two arguments, returning it in the third argument. **mtx()** provides the inverse of **xmtx()**. To release the storage allocated by **mtx()**, use **free()** (see **malloc(3V)**).

Use the **-lmp** loader option to obtain access to these functions.

#### DIAGNOSTICS

Illegal operations and running out of memory produce messages and core images.

#### FILES

**/usr/lib/libmp.a**

#### SEE ALSO

**malloc(3V)**



**NAME**

`msync` – synchronize memory with physical storage

**SYNOPSIS**

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

int msync(addr, len, flags)
caddr_t addr; size_t len; int flags;
```

**DESCRIPTION**

`msync()` writes all modified copies of pages over the range [*addr*, *addr + len*) to their permanent storage locations. `msync()` optionally invalidates any copies so that further references to the pages will be obtained by the system from their permanent storage locations.

Values for *flags* are defined in `<sys/mman.h>` as:

```
#define MS_ASYNC      0x1      /* Return immediately */
#define MS_INVALIDATE 0x2      /* Invalidate mappings */
```

and are used to control the behavior of `msync()`. One or more flags may be specified in a single call.

`MS_ASYNC` returns immediately once all I/O operations are scheduled; normally, `msync()` will not return until all I/O operations are complete. `MS_INVALIDATE` invalidates all cached copies of data from memory objects, requiring them to be re-obtained from the object's permanent storage location upon the next reference.

`msync()` should be used by programs that require a memory object to be in a known state, for example in building transaction facilities.

**RETURN VALUES**

`msync()` returns:

- 0        on success.
- 1       on failure and sets `errno` to indicate the error.

**ERRORS**

- EINVAL**        *addr* is not a multiple of the page size as returned by `getpagesize(2)`.  
*flags* is not some combination of `MS_ASYNC` or `MS_INVALIDATE`.
- EIO**            An I/O error occurred while reading from or writing to the file system.
- ENOMEM**        Addresses in the range [*addr*, *addr + len*) are outside the valid range for the address space of a process, or specify one or more pages that are not mapped.
- EPERM**        `MS_INVALIDATE` was specified and one or more of the pages is locked in memory.

**SEE ALSO**

`mctl(2)`, `mmap(2)`

## NAME

ndbm, dbm\_open, dbm\_close, dbm\_fetch, dbm\_store, dbm\_delete, dbm\_firstkey, dbm\_nextkey, dbm\_error, dbm\_clearerr – data base subroutines

## SYNOPSIS

```
#include <ndbm.h>

typedef struct {
    char *dptr;
    int dsize;
} datum;

DBM *dbm_open(file, flags, mode)
char *file;
int flags, mode;

void dbm_close (db)
DBM *db;

datum dbm_fetch(db, key)
DBM *db;
datum key;

int dbm_store(db, key, content, flags)
DBM *db;
datum key, content;
int flags;

int dbm_delete(db, key)
DBM *db;
datum key;

datum dbm_firstkey(db)
DBM *db;

datum dbm_nextkey(db)
DBM *db;

int dbm_error(db)
DBM *db;

int dbm_clearerr(db)
DBM *db;
```

## DESCRIPTION

These functions maintain key/content pairs in a data base. The functions will handle very large (a billion blocks) databases and will access a keyed item in one or two file system accesses. This package replaces the earlier **dbm(3X)** library, which managed only a single database.

*keys* and *contents* are described by the **datum** typedef. A **datum** specifies a string of *dsize* bytes pointed to by *dptr*. Arbitrary binary data, as well as normal ASCII strings, are allowed. The data base is stored in two files. One file is a directory containing a bit map and has **.dir** as its suffix. The second file contains all data and has **.pag** as its suffix.

Before a database can be accessed, it must be opened by **dbm\_open**. This will open and/or create the files *file.dir* and *file.pag* depending on the flags parameter (see **open(2V)**).

A database is closed by calling **dbm\_close**.

Once open, the data stored under a key is accessed by **dbm\_fetch()** and data is placed under a key by **dbm\_store**. The *flags* field can be either **DBM\_INSERT** or **DBM\_REPLACE**. **DBM\_INSERT** will only insert new entries into the database and will not change an existing entry with the same key. **DBM\_REPLACE** will replace an existing entry if it has the same key. A key (and its associated

contents) is deleted by **dbm\_delete**. A linear pass through all keys in a database may be made, in an (apparently) random order, by use of **dbm\_firstkey()** and **dbm\_nextkey**. **dbm\_firstkey()** will return the first key in the database. **dbm\_nextkey()** will return the next key in the database. This code will traverse the data base:

```
for (key = dbm_firstkey(db); key.dptr != NULL; key = dbm_nextkey(db))
```

**dbm\_error()** returns non-zero when an error has occurred reading or writing the database. **dbm\_clearerr()** resets the error condition on the named database.

#### SEE ALSO

**ar(1V)**, **cat(1V)**, **cp(1)**, **tar(1)**, **open(2V)**, **dbm(3X)**

#### DIAGNOSTICS

All functions that return an **int** indicate errors with negative values. A zero return indicates no error. Routines that return a **datum** indicate errors with a NULL (0) *dptr*. If **dbm\_store** called with a *flags* value of **DBM\_INSERT** finds an existing entry with the same key it returns 1.

#### BUGS

The **.pag** file will contain holes so that its apparent size is about four times its actual content. Older versions of the UNIX operating system may create real file blocks for these holes when touched. These files cannot be copied by normal means (**cp(1)**, **cat(1V)**, **tar(1)**, **ar(1V)**) without filling in the holes.

*dptr* pointers returned by these subroutines point into static storage that is changed by subsequent calls.

The sum of the sizes of a key/content pair must not exceed the internal block size (currently 4096 bytes). Moreover all key/content pairs that hash together must fit on a single block. **dbm\_store()** will return an error in the event that a disk block fills with inseparable data.

**dbm\_delete()** does not physically reclaim file space, although it does make it available for reuse.

The order of keys presented by **dbm\_firstkey()** and **dbm\_nextkey()** depends on a hashing function, not on anything interesting.

There are no interlocks and no reliable cache flushing; thus concurrent updating and reading is risky.

**NAME**

nice – change nice value of a process

**SYNOPSIS**

**int** nice(*incr*)

**DESCRIPTION**

The nice value of the process is changed by *incr*. Positive nice values get less service than normal. See **nice(1)** for a discussion of the relationship of nice value and scheduling priority.

A nice value of 10 is recommended to users who wish to execute long-running programs without undue impact on system performance.

Negative increments are illegal, except when specified by the super-user. The nice value is limited to the range -20 (most urgent) to 19 (least). Requests for values above or below these limits result in the nice value being set to the corresponding limit.

The nice value of a process is passed to a child process by **fork(2V)**. For a privileged process to return to normal nice value from an unknown state, **nice()** should be called successively with arguments -40 (goes to nice value -20 because of truncation), 20 (to get to 0), then 0 (to maintain compatibility with previous versions of this call).

**SYSTEM V DESCRIPTION**

The maximum allowed value for *incr* is 40 (least urgent).

**RETURN VALUES**

**nice()** returns:

0       on success.

-1       on failure and sets **errno** to indicate the error.

**SYSTEM V RETURN VALUES**

**nice()** returns the new nice value on success. On failure, it returns -1 and sets **errno** to indicate the error.

**ERRORS**

The nice value is not changed if:

**EACCES**       The value of *incr* specified was negative, and the effective user ID is not super-user.

**SYSTEM V ERRORS**

The nice value is not changed if:

**EPERM**       The value of *incr* specified was negative, or greater than 40, and the effective user ID is not super-user.

**SEE ALSO**

**nice(1)**, **fork(2V)**, **getpriority(2)**, **pstat(8)**, **renice(8)**

**NAME**

`nl_langinfo` – language information

**SYNOPSIS**

```
#include <nl_types.h>
#include <langinfo.h>

char *nl_langinfo(item)
nl_item item;
```

**DESCRIPTION**

`nl_langinfo()` returns a pointer to a null-terminated string containing information relevant to a particular language or cultural area defined in the program's locale. The manifest constant names and values of *item* are defined in `<langinfo.h>`. For example:

```
nl_langinfo(ABDAY_1);
```

would return a pointer to the string 'Dom' if the identified language was Portuguese, and 'Sun' if the identified language was English.

**RETURN VALUES**

In a locale where *langinfo* data is not defined, `nl_langinfo()` returns a pointer to the corresponding string in the "C" locale. In all locales `nl_langinfo()` returns a pointer to an empty string if *item* contains an invalid setting.

**SEE ALSO**

`setlocale(3V)`, `environ(5V)`

**NAME**

nlist – get entries from symbol table

**SYNOPSIS**

```
#include <nlist.h>
```

```
int nlist(filename, nl)
```

```
char *filename;
```

```
struct nlist *nl;
```

**DESCRIPTION**

**nlist()** examines the symbol table from the executable image whose name is pointed to by *filename*, and selectively extracts a list of values and puts them in the array of **nlist()** structures pointed to by *nl*. The name list pointed to by *nl* consists of an array of structures containing names, types and values. The *n\_name* field of each such structure is taken to be a pointer to a character string representing a symbol name. The list is terminated by an entry with a NULL pointer (or a pointer to a null string) in the *n\_name* field. For each entry in *nl*, if the named symbol is present in the executable image's symbol table, its value and type are placed in the *n\_value* and *n\_type* fields. If a symbol cannot be located, the corresponding *n\_type* field of *nl* is set to zero.

**RETURN VALUES**

On success, **nlist()** returns the number of symbols that were not located in the symbol table. On failure, it returns -1 and sets all of the *n\_type* fields in members of the array pointed to by *nl* to zero.

**SYSTEM V RETURN VALUES**

**nlist()** returns 0 on success.

**SEE ALSO**

**a.out(5)**, **coff(5)**

**NOTES**

On Sun-2, Sun-3, and Sun-4 systems, type entries are set to 0 if the file cannot be read or if it does not contain a valid name list.

On Sun386i systems, the type entries may be zero even when the name list succeeded, but the value entries will be zero only when the file cannot be read or does not contain a valid name list. Therefore, on Sun386i systems, the value entry can be used to determine whether the command succeeded.

**NAME**

`on_exit` – name termination handler

**SYNOPSIS**

```
int on_exit(procp, arg)
void (*procp)();
caddr_t arg;
```

**DESCRIPTION**

`on_exit()` names a routine to be called after a program calls `exit(3)` or returns normally, and before its process terminates. The routine named is called as

```
(*procp)(status, arg);
```

where *status* is the argument with which `exit()` was called, or zero if *main* returns. Typically, *arg* is the address of an argument vector to *(\*procp)*, but may be an integer value. Several calls may be made to `on_exit`, specifying several termination handlers. The order in which they are called is the reverse of that in which they were given to `on_exit`.

**SEE ALSO**

`gprof(1)`, `tcov(1)`, `exit(3)`

**DIAGNOSTICS**

`on_exit()` returns zero normally, or nonzero if the procedure name could not be stored.

**NOTES**

This call is specific to the SunOS operating system and should not be used if portability is a concern. Standard I/O exit processing is always done last.

**NAME**

pause – stop until signal

**SYNOPSIS**

**int** pause()

**DESCRIPTION**

**pause()** never returns normally. It is used to give up control while waiting for a signal from **kill(2V)** or an interval timer, see **getitimer(2)**. Upon termination of a signal handler started during a **pause**, **pause()** will return.

**RETURN VALUES**

When it returns, **pause()** returns **-1**.

**ERRORS**

When it returns, **pause()** sets **errno** to:

**EINTR**            A signal is caught by the calling process and control is returned from the signal-catching function.

**SEE ALSO**

**kill(2V)**, **getitimer(2)**, **select(2)**, **sigpause(2V)**



**NAME**

**perror, errno** – system error messages

**SYNOPSIS**

```
void perror(s)  
char *s;  
#include <errno.h>  
int sys_nerr;  
char *sys_errlist[ ];  
int errno;
```

**DESCRIPTION**

**perror()** produces a short error message on the standard error describing the last error encountered during a call to a system or library function. If *s* is not a NULL pointer and does not point to a null string, the string it points to is printed, followed by a colon, followed by a space, followed by the message and a NEWLINE. If *s* is a NULL pointer or points to a null string, just the message is printed, followed by a NEWLINE. To be of most use, the argument string should include the name of the program that incurred the error. The error number is taken from the external variable **errno** (see **intro(2)**), which is set when errors occur but not cleared when non-erroneous calls are made.

To simplify variant formatting of messages, the vector of message strings **sys\_errlist** is provided; **errno** can be used as an index in this table to get the message string without the newline. **sys\_nerr** is the number of messages provided for in the table; it should be checked because new error codes may be added to the system before they are added to the table.

**SEE ALSO**

**intro(2), psignal(3)**

**NAME**

**plock** – lock process, text, or data segment in memory

**SYNOPSIS**

```
#include <sys/lock.h>
```

```
int plock(op)
```

```
int op;
```

**DESCRIPTION**

**plock()** allows the calling process to lock its text segment (text lock), its data segment (data lock), or both its text and data segments (process lock) into memory. Locked segments are immune to all routine swapping. **plock()** also allows these segments to be unlocked. The effective user ID of the calling process must be super-user to use this call. *op* specifies the following:

<b>PROCLOCK</b>	lock text and data segments into memory (process lock)
<b>TXTLOCK</b>	lock text segment into memory (text lock)
<b>DATLOCK</b>	lock data segment into memory (data lock)
<b>UNLOCK</b>	remove locks

**RETURN VALUES**

**plock()** returns:

0 on success.

-1 on failure and sets **errno** to indicate the error.

**ERRORS**

**EAGAIN** Not enough memory.

**EINVAL** *op* is equal to **PROCLOCK** and a process lock, a text lock, or a data lock already exists on the calling process.

*op* is equal to **TXTLOCK** and a text lock, or a process lock already exists on the calling process.

*op* is equal to **DATLOCK** and a data lock, or a process lock already exists on the calling process.

*op* is equal to **UNLOCK** and no type of lock exists on the calling process.

**EPERM** The effective user ID of the calling process is not super-user.

**SEE ALSO**

**execve(2V)**, **exit(2V)**, **fork(2V)**

## NAME

plot, openpl, erase, label, line, circle, arc, move, cont, point, linemod, space, closepl – graphics interface

## SYNOPSIS

```

openpl()
erase()
label(s)
char s[ ];
line(x1, y1, x2, y2)
circle(x, y, r)
arc(x, y, x0, y0, x1, y1)
move(x, y)
cont(x, y)
point(x, y)
linemod(s)
char s[ ];
space(x0, y0, x1, y1)
closepl()

```

## AVAILABILITY

These routines are available with the *Graphics* software installation option. Refer to *Installing SunOS 4.1* for information on how to install optional software.

## DESCRIPTION

LP These subroutines generate graphic output in a relatively device-independent manner. See **plot(5)** for a description of their effect. **openpl()** must be used before any of the others to open the device for writing. **closepl()** flushes the output.

String arguments to **label()** and **linemod()** are null-terminated and do not contain NEWLINE characters.

Various flavors of these functions exist for different output devices. They are obtained by the following **ld(1)** options:

<b>-lplot</b>	device-independent graphics stream on standard output for <b>plot(1G)</b> filters
<b>-l300</b>	GSI 300 terminal
<b>-l300s</b>	GSI 300S terminal
<b>-l450</b>	GSI 450 terminal
<b>-l4014</b>	Tektronix 4014 terminal
<b>-lplotaed</b>	AED 512 color graphics terminal
<b>-lplotbg</b>	BBN bitgraph graphics terminal
<b>-lplotdumb</b>	Dumb terminals without cursor addressing or line printers
<b>-lplotgigi</b>	DEC Gigi terminals
<b>-lplot2648</b>	Hewlett Packard 2648 graphics terminal
<b>-lplot7221</b>	Hewlett Packard 7221 graphics terminal
<b>-lplotimagen</b>	Imagen laser printer (default 240 dots-per-inch resolution).

**FILES**

**/usr/lib/libplot.a**  
**/usr/lib/lib300.a**  
**/usr/lib/lib300s.a**  
**/usr/lib/lib450.a**  
**/usr/lib/lib4014.a**  
**/usr/lib/libplotaed.a**  
**/usr/lib/libplotbg.a**  
**/usr/lib/libplotdumb.a**  
**/usr/lib/libplotgigi.a**  
**/usr/lib/libplot2648.a**  
**/usr/lib/libplot7221.a**  
**/usr/lib/libplotimagen.a**

**SEE ALSO**

**graph(1G), ld(1), plot(1G), plot(5)**

**NAME**

**popen, pclose** – open or close a pipe (for I/O) from or to a process

**SYNOPSIS**

```
#include <stdio.h>

FILE *popen(command, type)
char *command, *type;

pclose(stream)
FILE *stream;
```

**DESCRIPTION**

The arguments to **popen()** are pointers to null-terminated strings containing, respectively, a shell command line and an I/O mode, either **r** for reading or **w** for writing. **popen()** creates a pipe between the calling process and the command to be executed. The value returned is a stream pointer such that one can write to the standard input of the command, if the I/O mode is **w**, by writing to the file stream; and one can read from the standard output of the command, if the I/O mode is **r**, by reading from the file stream.

A stream opened by **popen()** should be closed by **pclose()**, which waits for the associated process to terminate and returns the exit status of the command.

Because open files are shared, a type **r** command may be used as an input filter, reading its standard input (which is also the standard output of the process doing the **popen()**) and providing filtered input on the stream, and a type **w** command may be used as an output filter, reading a stream of output written to the stream process doing the **popen()** and further filtering it and writing it to its standard output (which is also the standard input of the process doing the **popen()**).

**popen()** always calls **sh(1)**, never **csh(1)**.

**SEE ALSO**

**csh(1)**, **sh(1)**, **pipe(2V)**, **wait(2V)**, **fclose(3V)**, **fopen(3V)**, **system(3)**

**DIAGNOSTICS**

**popen()** returns a NULL pointer if the pipe or process cannot be created, or if it cannot allocate as much memory as it needs.

**pclose()** returns **-1** if stream is not associated with a 'popened' command.

**BUGS**

If the original and 'popened' processes concurrently read or write a common file, neither should use buffered I/O, because the buffering gets all mixed up. Similar problems with an output filter may be forestalled by careful buffer flushing, for instance, with **fflush()**; see **fclose(3V)**.

## NAME

`pmap_getmaps`, `pmap_getport`, `pmap_rmtcall`, `pmap_set`, `pmap_unset`, `xdr_pamp`, `xdr_pmaplist` – library routines for RPC bind service

## DESCRIPTION

These routines allow client C programs to make procedure calls to the RPC binder service. `portmap(1)` maintains a list of mappings between programs and their universal addresses.

## Routines

```
#include <rpc/rpc.h>
```

```
struct pmaplist * pmap_getmaps(addr)
struct sockaddr_in *addr;
```

Return a list of the current RPC program-to-address mappings on the host located at IP address *\*addr*. This routine returns NULL if the remote `portmap` service could not be contacted. The command `'rpcinfo -p'` uses this routine (see `rpcinfo(8C)`).

```
u_short pmap_getport(addr, prognum, versnum, protocol)
struct sockaddr_in *addr;
u_long prognum, versnum, protocol;
```

Return the port number on which waits a service that supports program number *prognum*, version *versnum*, and speaks the transport protocol *protocol*. The address is returned in *addr*, which should be preallocated. The value of *protocol* can be either `IPPROTO_UDP` or `IPPROTO_TCP`. A return value of zero means that the mapping does not exist or that the RPC system failed to contact the remote `portmap` service. In the latter case, the global variable `rpc_createer` (see `rpc_clnt_create(3N)`) contains the RPC status. If the requested version number is not registered, but at least a version number is registered for the given program number, the call returns a port number. Note: `pmap_getport()` returns the port number in host byte order. Some other network routines may require the port number in network byte order. For example, if the port number is used as part of the `sockaddr_in` structure, then it should be converted to network byte order using `htons(3N)`.

```
enum clnt_stat pmap_rmtcall(addr, prognum, versnum, procnum, inproc, in, outproc, out, timeout, portp)
struct sockaddr_in *addr;
u_long prognum, versnum, procnum;
char *in, *out;
xdrproc_t inproc, outproc;
struct timeval timeout;
u_long *portp;
```

Request that the `portmap` on the host at IP address *\*addr* make an RPC on the behalf of the caller to a procedure on that host. *\*portp* is modified to the program's port number if the procedure succeeds. The definitions of other parameters are discussed in `callrpc()` and `clnt_call()` (see `rpc_clnt_calls(3N)`).

Warning: If the requested remote procedure is not registered with the remote `portmap` then no error response is returned and the call times out. Also, no authentication is done.

```
bool_t pmap_set(prognum, versnum, protocol, port)
u_long prognum, versnum;
int protocol;
u_short port;
```

Registers a mapping between the triple [*prognum,versnum.protocol*] and *port* on the local machine's `portmap` service. The value of *protocol* can be either `IPPROTO_UDP` or `IPPROTO_TCP`. This routine returns TRUE if it succeeds, FALSE otherwise. It is called by servers to register themselves with the local `portmap`. Automatically done by `svc_register()`.

**bool\_t pmap\_unset(prognum, versnum)**  
**u\_long prognum, versnum;**

Deregisters all mappings between the triple [*prognum,versnum,\**] and ports on the local machine's **portmap** service. It is called by servers to deregister themselves with the local **portmap**. This routine returns TRUE if it succeeds, FALSE otherwise.

**bool\_t xdr\_pmap(xdrs, regp)**  
**XDR \*xdrs;**  
**struct pmap \*regp;**

Used for creating parameters to various **portmap** procedures, externally. This routine is useful for users who wish to generate these parameters without using the **pmap** interface. This routine returns TRUE if it succeeds, FALSE otherwise.

**bool\_t xdr\_pmaplist(xdrs, rp)**  
**XDR \*xdrs;**  
**struct pmaplist \*\*rp;**

Used for creating a list of port mappings, externally. This routine is useful for users who wish to generate these parameters without using the **pmap** interface. This routine returns TRUE if it succeeds, FALSE otherwise.

**SEE ALSO**

**rpc(3N), portmap(8C), rpcinfo(8C)**

## NAME

printf, fprintf, sprintf – formatted output conversion

## SYNOPSIS

```
#include <stdio.h>

int printf(format [ , arg... ])
char *format;

int fprintf(stream, format [ , arg... ])
FILE *stream;
char *format;

char *sprintf(s, format [ , arg... ])
char *s, *format;
```

## SYSTEM V SYNOPSIS

The routines above are available as shown, except:

```
int sprintf(s, format [ , arg... ])
char *s, *format;
```

The following are provided for XPG2 compatibility:

```
#define nl_printf      printf
#define nl_fprintf    fprintf
#define nl_sprintf    sprintf
```

## DESCRIPTION

**printf()** places output on the standard output stream **stdout**. **fprintf()** places output on the named output **stream**. **sprintf()** places “output”, followed by the null character (**\0**), in consecutive bytes starting at **\*s**; it is the user’s responsibility to ensure that enough storage is available.

Each of these functions converts, formats, and prints its *args* under control of the *format*. The *format* is a character string which contains two types of objects: plain characters, which are simply copied to the output stream, and conversion specifications, each of which causes conversion and printing of zero or more *args*. The results are undefined if there are insufficient *args* for the format. If the format is exhausted while *args* remain, the excess *args* are simply ignored.

Each conversion specification is introduced by either the **%** character or by the character sequence **%digit\$**, after which the following appear in sequence:

- Zero or more *flags*, which modify the meaning of the conversion specification.
- An optional decimal digit string specifying a minimum *field width*. If the converted value has fewer characters than the field width, it will be padded on the left (or right, if the left-adjustment flag ‘-’, described below, has been given) to the field width. The padding is with blanks unless the field width digit string starts with a zero, in which case the padding is with zeros.
- A *precision* that gives the minimum number of digits to appear for the **d**, **i**, **o**, **u**, **x**, or **X** conversions, the number of digits to appear after the decimal point for the **e**, **E**, and **f** conversions, the maximum number of significant digits for the **g** and **G** conversion, or the maximum number of characters to be printed from a string in **s** conversion. The precision takes the form of a period (.) followed by a decimal digit string; a null digit string is treated as zero. Padding specified by the precision overrides the padding specified by the field width.
- An optional **l** (ell) specifying that a following **d**, **i**, **o**, **u**, **x**, or **X** conversion character applies to a long integer *arg*. An **l** before any other conversion character is ignored.
- A character that indicates the type of conversion to be applied.



A field width or precision or both may be indicated by an asterisk (\*) instead of a digit string. In this case, an integer *arg* supplies the field width or precision. The *arg* that is actually converted is not fetched until the conversion letter is seen, so the *args* specifying field width or precision must appear *before* the *arg* (if any) to be converted. A negative field width argument is taken as a ‘-’ flag followed by a positive field width. If the precision argument is negative, it will be changed to zero.

The flag characters and their meanings are:

- The result of the conversion will be left-justified within the field.
- + The result of a signed conversion will always begin with a sign (+ or -).
- blank If the first character of a signed conversion is not a sign, a blank will be prefixed to the result. This implies that if the blank and + flags both appear, the blank flag will be ignored.
- # This flag specifies that the value is to be converted to an “alternate form”. For *c*, *d*, *i*, *s*, and *u* conversions, the flag has no effect. For *o* conversion, it increases the precision to force the first digit of the result to be a zero. For *x* or *X* conversion, a non-zero result will have *0x* or *0X* prefixed to it. For *e*, *E*, *f*, *g*, and *G* conversions, the result will always contain a decimal point, even if no digits follow the point (normally, a decimal point appears in the result of these conversions only if a digit follows it). For *g* and *G* conversions, trailing zeroes will *not* be removed from the result (which they normally are).

The conversion characters and their meanings are:

#### **d,i,o,p,u,x,X**

The integer *arg* is converted to signed decimal (*d* or *i*), unsigned octal (*o*), unsigned decimal (*u*), or unsigned hexadecimal notation (*x*, *p*, and *X*), respectively; the letters *abcdef* are used for *x* and *p* conversion and the letters *ABCDEF* for *X* conversion. The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeroes. For compatibility with older versions, padding with leading zeroes may alternatively be specified by prepending a zero to the field width. This does not imply an octal value for the field width. The default precision is 1. The result of converting a zero value with a precision of zero is a null string.

- f** The float or double *arg* is converted to decimal notation in the style “[–]ddd.ddd” where the number of digits after the decimal point is equal to the precision specification. If the precision is missing, 6 digits are given; if the precision is explicitly 0, no digits and no decimal point are printed.
- e,E** The float or double *arg* is converted in the style “[–]d.ddde±ddd,” where there is one digit before the decimal point and the number of digits after it is equal to the precision; when the precision is missing, 6 digits are produced; if the precision is zero, no decimal point appears. The *E* format code will produce a number with *E* instead of *e* introducing the exponent. The exponent always contains at least two digits.
- g,G** The float or double *arg* is printed in style *f* or *e* (or in style *E* in the case of a *G* format code), with the precision specifying the number of significant digits. The style used depends on the value converted: style *e* or *E* will be used only if the exponent resulting from the conversion is less than -4 or greater than the precision. Trailing zeroes are removed from the result; a decimal point appears only if it is followed by a digit.

The *e*, *E*, *f*, *g*, and *G* formats print IEEE indeterminate values (infinity or not-a-number) as “Infinity” or “NaN” respectively.

- c** The character *arg* is printed.
- s** The *arg* is taken to be a string (character pointer) and characters from the string are printed until a null character ( $\backslash 0$ ) is encountered or until the number of characters indicated by the precision specification is reached. If the precision is missing, it is taken to be infinite, so all characters up to the first null character are printed. A NULL value for *arg* will yield undefined results.

**n** The argument *arg* is a pointer to an integer into which is written the number of characters written to the output so far by this call to one of the `printf()` functions. No argument is converted.

**%** Print a `%`; no argument is converted.

In no case does a non-existent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is simply expanded to contain the conversion result. Padding takes place only if the specified field width exceeds the actual width. Characters generated by `printf()` and `fprintf()` are printed as if `putc(3S)` had been called.

All forms of the `printf()` functions allow for the insertion of a language dependent radix character in the output string. The radix character is defined by the program's locale (category `LC_NUMERIC`). In the "C" locale, or in a locale where the radix character is not defined, the radix character defaults to `'.'`.

Conversions can be applied to the *n*th argument in the argument list, rather than the next unused argument. In this case, the conversion character `%` is replaced by the sequence `%digit$`, where *digit* is a decimal integer *n* in the range [1,9], giving the position of the argument in the argument list. This feature provides for the definition of format strings that select arguments in an order appropriate to specific languages.

In format strings containing the `%digit$` form of a conversion specification, a field width or precision may be indicated by the sequence `*digit$`, where *digit* is a decimal integer in the range [1,9] giving the position in the argument list of an integer *arg* containing the field width or precision.

The format string can contain either numbered argument specifications (that is, `%digit$` and `*digit$`), or unnumbered argument specifications (that is `%` and `*`), but not both. The results of mixing numbered and unnumbered specifications is undefined. When numbered argument specifications are used, specifying the *n*th argument requires that all the leading arguments, from the first to the (*n*-1)th be specified in the format string.

#### SYSTEM V DESCRIPTION

XPG2 requires that `nl_printf`, `nl_fprintf` and `nl_sprintf` be defined as `printf`, `fprintf` and `sprintf`, respectively for backward compatibility

#### RETURN VALUES

On success, `printf()` and `fprintf()` return the number of characters transmitted, excluding the null character. On failure, they return EOF.

`sprintf()` returns *s*.

#### SYSTEM V RETURN VALUES

On success, `sprintf()` returns the number of characters transmitted, excluding the null character. On failure, it returns EOF.

#### EXAMPLES

```
printf(format, weekday, month, day, hour, min);
```

In American usage, *format* could be a pointer to the string:

```
"%s, %s %d, %d:%.2d\n"
```

producing the message:

```
Sunday, July 3,10:02
```

Whereas for German usage, *format* could be a pointer to the string:

```
"%1$s, %3$d.%2$s, %4$d:%5$.2d\n"
```

producing the message:

```
Sonntag, 3.Juli,10:02
```

To print  $\pi$  to 5 decimal places:

```
printf("pi = %.5f", 4 * atan(1.0));
```

**SEE ALSO**

**econvert(3), putc(3S), scanf(3V), setlocale(3V), varargs(3), vprintf(3V)**

**BUGS**

Very wide fields (>128 characters) fail.

**NAME**

prof – profile within a function

**SYNOPSIS**

```
#define MARK
#include <prof.h>

void MARK (name)
```

**DESCRIPTION**

MARK introduces a mark called *name* that is treated the same as a function entry point. Execution of the mark adds to a counter for that mark, and program-counter time spent is accounted to the immediately preceding mark or to the function if there are no preceding marks within the active function.

*name* may be any combination of up to six letters, numbers or underscores. Each *name* in a single compilation must be unique, but may be the same as any ordinary program symbol.

For marks to be effective, the symbol MARK must be defined before the header file <prof.h> is included. This may be defined by a preprocessor directive as in the synopsis, or by a command line argument, such as:

```
cc -p -DMARK foo.c
```

If MARK is not defined, the MARK (*name*) statements may be left in the source files containing them and will be ignored.

**EXAMPLE**

In this example, marks can be used to determine how much time is spent in each loop. Unless this example is compiled with MARK defined on the command line, the marks are ignored.

```
#include <prof.h>
func( )
{
    int i, j;
    .
    .
    .
    MARK (loop1);
    for (i = 0; i < 2000; i++) {
        ...
    }
    MARK (loop2);
    for (j = 0; j < 2000; j++) {
        ...
    }
}
```

**SEE ALSO**

prof(1), profil(2), monitor(3)

**NAME**

**psignal, sys\_siglist** – system signal messages

**SYNOPSIS**

```
psignal(sig, s)  
unsigned sig;  
char *s;  
  
char *sys_siglist[ ];
```

**DESCRIPTION**

**psignal()** produces a short message on the standard error file describing the indicated signal. First the argument string *s* is printed, then a colon, then the name of the signal and a NEWLINE. Most usefully, the argument string is the name of the program which incurred the signal. The signal number should be from among those found in **<signal.h>**.

To simplify variant formatting of signal names, the vector of message strings **sys\_siglist()** is provided; the signal number can be used as an index in this table to get the signal name without the newline. The define **NSIG** defined in **<signal.h>** is the number of messages provided for in the table; it should be checked because new signals may be added to the system before they are added to the table.

**SEE ALSO**

**perror(3), signal(3V)**

## NAME

`putc`, `putchar`, `fputc`, `putw` – put character or word on a stream

## SYNOPSIS

```
#include <stdio.h>

int putc(c, stream)
char c;
FILE *stream;

int putchar(c)
char c;

int fputc(c, stream)
char c;
FILE *stream;

int putw(w, stream)
int w;
FILE *stream;
```

## DESCRIPTION

`putc()` writes the character *c* onto the standard I/O output stream *stream* (at the position where the file pointer, if defined, is pointing). It returns the character written.

`putchar(c)` is defined as `putc(c, stdout)`. `putc()` and `putchar()` are macros.

`fputc()` behaves like `putc()`, but is a function rather than a macro. `fputc()` runs more slowly than `putc()`, but it takes less space per invocation and its name can be passed as an argument to a function.

`putw()` writes the C `int` (word) *w* to the standard I/O output stream *stream* (at the position of the file pointer, if defined). The size of a word is the size of an integer and varies from machine to machine. `putw()` neither assumes nor causes special alignment in the file.

Output streams are by default buffered if the output refers to a file and line-buffered if the output refers to a terminal. When an output stream is unbuffered, information is queued for writing on the destination file or terminal as soon as written; when it is buffered, many characters are saved up and written as a block. When it is line-buffered, each line of output is queued for writing on the destination terminal as soon as the line is completed (that is, as soon as a `NEWLINE` character is written or terminal input is requested). `setbuf(3V)`, `setbuffer()`, or `setvbuf()` may be used to change the stream's buffering strategy.

## SEE ALSO

`fclose(3V)`, `ferror(3V)`, `fopen(3V)`, `fread(3S)`, `getc(3V)`, `printf(3V)`, `puts(3S)`, `setbuf(3V)`

## DIAGNOSTICS

On success, `putc()`, `fputc()`, and `putchar()` return the value that was written. On error, those functions return the constant `EOF`. `putw()` returns `ferror(stream)`, so that it returns 0 on success and 1 on failure.

## BUGS

Because it is implemented as a macro, `putc()` treats a *stream* argument with side effects improperly. In particular, `putc(c, *f++)`; does not work sensibly. `fputc()` should be used instead.

Errors can occur long after the call to `putc()`.

Because of possible differences in word length and byte ordering, files written using `putw()` are machine-dependent, and may not be read using `getw()` on a different processor.

**NAME**

putenv – change or add value to environment

**SYNOPSIS**

```
int putenv(string)
char *string;
```

**DESCRIPTION**

*string* points to a string of the form '*name=value*'. **putenv()** makes the value of the environment variable *name* equal to *value* by altering an existing variable or creating a new one. In either case, the string pointed to by *string* becomes part of the environment, so altering the string will change the environment. The space used by *string* is no longer used once a new string-defining *name* is passed to **putenv()**.

**SEE ALSO**

**execve(2V)**, **getenv(3V)**, **malloc(3V)**, **environ(5V)**

**DIAGNOSTICS**

**putenv()** returns non-zero if it was unable to obtain enough space using **malloc(3V)** for an expanded environment, otherwise zero.

**WARNINGS**

**putenv()** manipulates the environment pointed to by *environ*, and can be used in conjunction with **getenv()**. However, *envp* (the third argument to *main*) is not changed.

This routine uses **malloc(3V)** to enlarge the environment.

After **putenv()** is called, environmental variables are not in alphabetical order.

A potential error is to call **putenv()** with an automatic variable as the argument, then exit the calling function while *string* is still part of the environment.

**NAME**

putpwent – write password file entry

**SYNOPSIS**

```
#include <pwd.h>

int putpwent(p, f)
struct passwd *p;
FILE *f;
```

**DESCRIPTION**

**putpwent()** is the inverse of **getpwent(3V)**. Given a pointer to a **passwd** structure created by **getpwent()** (or **getpwuid()** or **getpwnam()**), **putpwent()** writes a line on the stream *f*, which matches the format of lines in the password file */etc/passwd*.

**FILES**

*/etc/passwd*

**SEE ALSO**

**getpwent(3V)**

**DIAGNOSTICS**

**putpwent()** returns non-zero if an error was detected during its operation, otherwise zero.

**WARNING**

The above routine uses **<stdio.h>**, which increases the size of programs, not otherwise using standard I/O, more than might be expected.

**BUGS**

This routine is of limited utility, since most password files are maintained as Network Information Service (NIS) files, and cannot be updated with this routine.

**NOTES**

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed.



**NAME**

puts, fputs – put a string on a stream

**SYNOPSIS**

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

```
puts(s)
```

```
char *s;
```

```
fputs(s, stream)
```

```
char *s;
```

```
FILE *stream;
```

**DESCRIPTION**

**puts()** writes the null-terminated string pointed to by *s*, followed by a NEWLINE character, to the standard output stream **stdout**.

**fputs()** writes the null-terminated string pointed to by *s* to the named output stream.

Neither function writes the terminal null character.

**DIAGNOSTICS**

Both routines return EOF on error. This will happen if the routines try to write on a file that has not been opened for writing.

**NOTES**

**puts()** appends a NEWLINE while **fputs()** does not.

**SEE ALSO**

**ferror(3V)**, **fopen(3V)**, **fread(3S)**, **printf(3V)**, **putc(3S)**

**NAME**

**pwdauth, grpauth** – password authentication routines

**SYNOPSIS**

```
int pwdauth(user, password)
```

```
char *user;
```

```
char *password;
```

```
int grpauth(group, password)
```

```
char *group;
```

```
char *password;
```

**DESCRIPTION**

**pwdauth()** and **grpauth()** determine whether the given guess at a *password* is valid for the given *user* or *group*. If the *password* is valid, the functions return 0.

A *password* is valid if the password when encrypted matches the encrypted password in the appropriate file. For **pwdauth()**, if the **password.adjunct** file exists, the encrypted password will be in either the local or the Network Information Service (NIS) version of that file. Otherwise, either the local or NIS **passwd** file will be used. For **grpauth()**, the **group.adjunct** file (if it exists) or the **group** file (otherwise) will be checked on the local machine and then using the NIS service. In all cases, the local files will be checked before the NIS files. Also, if the adjunct files exist, the main file will never be used for authentication even if they include encrypted passwords.

Both **pwdauth()** and **grpauth()** interface to the authentication daemon, **rpc.pwdauthd**, to do the checking of the adjunct files. This daemon must be running on any system that provides password authentication.

**FILES**

**/etc/passwd**

**/etc/group**

**SEE ALSO**

**getgraent(3)**, **getgrent(3V)**, **getpwaent(3)**, **getpwent(3V)**, **pwdauthd(8C)**

**NOTES**

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed.

**NAME**

qsort – quicker sort

**SYNOPSIS**

```
qsort(base, nel, width, compar)
char *base;
int (*compar)();
```

**DESCRIPTION**

**qsort()** is an implementation of the quicker-sort algorithm. It sorts a table of data in place.

*base* points to the element at the base of the table. *nel* is the number of elements in the table. *width* is the size, in bytes, of each element in the table. *compar* is the name of the comparison function, which is called with two arguments that point to the elements being compared. As the function must return an integer less than, equal to, or greater than zero, so must the first argument to be considered be less than, equal to, or greater than the second.

**NOTES**

The pointer to the base of the table should be of type pointer-to-element, and cast to type pointer-to-character.

The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

The order in the output of two items which compare as equal is unpredictable.

**SEE ALSO**

**sort(1V)**, **bsearch(3)**, **lsearch(3)**, **string(3)**

**EXAMPLE**

The following program sorts a simple array:

```
static int intcompare(i,j)
int *i, *j;
{
    return(*i - *j);
}

main()
{
    int a[10];
    int i;

    a[0] = 9;
    a[1] = 8;
    a[2] = 7;
    a[3] = 6;
    a[4] = 5;
    a[5] = 4;
    a[6] = 3;
    a[7] = 2;
    a[8] = 1;
    a[9] = 0;

    qsort(a,10,sizeof(int),intcompare)
    for (i=0; i<10; i++) printf(" %d",a[i]);
    printf("\n");
}
```

**NAME**

rand, srand – simple random number generator

**SYNOPSIS**

**srand(seed)**

**int seed;**

**rand()**

**DESCRIPTION**

**rand()** uses a multiplicative congruential random number generator with period  $2^{32}$  to return successive pseudo-random numbers in the range from 0 to  $2^{31}-1$ .

**srand()** can be called at any time to reset the random-number generator to a random starting point. The generator is initially seeded with a value of 1.

**SYSTEM V DESCRIPTION**

**rand()** returns successive pseudo-random numbers in the range from 0 to  $2^{15}-1$ .

**SEE ALSO**

**drand48(3)**, **random(3)**

**NOTES**

The spectral properties of **rand()** leave a great deal to be desired. **drand48(3)** and **random(3)** provide much better, though more elaborate, random-number generators.

**BUGS**

The low bits of the numbers generated are not very random; use the middle bits. In particular the lowest bit alternates between 0 and 1.

## NAME

random, srand, initstate, setstate – better random number generator; routines for changing generators

## SYNOPSIS

```
long random()
srand(seed)
int seed;

char *initstate(seed, state, n)
unsigned seed;
char *state;
int n;

char *setstate(state)
char *state;
```

## DESCRIPTION

**random()** uses a non-linear additive feedback random number generator employing a default table of size 31 long integers to return successive pseudo-random numbers in the range from 0 to  $2^{31}-1$ . The period of this random number generator is very large, approximately  $16 \times (2^{31}-1)$ .

**random/srand** have (almost) the same calling sequence and initialization properties as **rand/srand**. The difference is that **rand(3V)** produces a much less random sequence — in fact, the low dozen bits generated by **rand** go through a cyclic pattern. All the bits generated by **random()** are usable. For example,

```
random()&01
```

will produce a random binary value.

Unlike **srand**, **srandom()** does not return the old seed; the reason for this is that the amount of state information used is much more than a single word. (Two other routines are provided to deal with restarting/changing random number generators). Like **rand(3V)**, however, **random()** will by default produce a sequence of numbers that can be duplicated by calling **srandom()** with *l* as the seed.

The **initstate()** routine allows a state array, passed in as an argument, to be initialized for future use. The size of the state array (in bytes) is used by **initstate()** to decide how sophisticated a random number generator it should use — the more state, the better the random numbers will be. (Current “optimal” values for the amount of state information are 8, 32, 64, 128, and 256 bytes; other amounts will be rounded down to the nearest known amount. Using less than 8 bytes will cause an error). The seed for the initialization (which specifies a starting point for the random number sequence, and provides for restarting at the same point) is also an argument. **initstate()** returns a pointer to the previous state information array.

Once a state has been initialized, the **setstate()** routine provides for rapid switching between states. **setstate()** returns a pointer to the previous state array; its argument state array is used for further random number generation until the next call to **initstate()** or **setstate()**.

Once a state array has been initialized, it may be restarted at a different point either by calling **initstate()** (with the desired seed, the state array, and its size) or by calling both **setstate()** (with the state array) and **srandom()** (with the desired seed). The advantage of calling both **setstate()** and **srandom()** is that the size of the state array does not have to be remembered after it is initialized.

With 256 bytes of state information, the period of the random number generator is greater than  $2^{69}$ , which should be sufficient for most purposes.

## SEE ALSO

**rand(3V)**

## EXAMPLES

```

/* Initialize and array and pass it in to initstate. */
static long state1[32] = {
    3,
    0x9a319039, 0x32d9c024, 0x9b663182, 0x5da1f342,
    0x7449e56b, 0xbcb1dbb0, 0xab5c5918, 0x946554fd,
    0x8c2e680f, 0xeb3d799f, 0xb11ee0b7, 0x2d436b86,
    0xda672e2a, 0x1588ca88, 0xe369735d, 0x904f35f7,
    0xd7158fd6, 0x6fa6f051, 0x616e6b96, 0xac94efdc,
    0xde3b81e0, 0xdf0a6fb5, 0xf103bc02, 0x48f340fb,
    0x36413f93, 0xc622c298, 0xf5a42ab8, 0x8a88d77b,
    0xf5ad9d0e, 0x8999220b, 0x27fb47b9
};

main()
{
    unsigned seed;
    int n;

    seed = 1;
    n = 128;
    initstate(seed, (char *) state1, n);

    setstate(state1);
    printf(" %d\n", random());
}

```

## DIAGNOSTICS

If `initstate()` is called with less than 8 bytes of state information, or if `setstate()` detects that the state information has been garbled, error messages are printed on the standard error output.

## WARNINGS

`initstate()` casts `state` to `(long *)`, so `state` must be long-aligned. If it is not long-aligned, on some architectures the program will dump core.

## BUGS

`random()` is only 2/3 as fast as `rand(3V)`.

## NAME

`rcmd`, `rresvport`, `ruserok` – routines for returning a stream to a remote command

## SYNOPSIS

```
int rcmd(ahost, inport, locuser, remuser, cmd, fd2p)
char **ahost;
unsigned short inport;
char *locuser, *remuser, *cmd;
int *fd2p

int rresvport(port)
int *port;

ruserok(rhost, super-user, ruser, luser)
char *rhost;
int super-user;
char *ruser, *luser;
```

## DESCRIPTION

`rcmd()` is a routine used by the super-user to execute a command on a remote machine using an authentication scheme based on reserved port numbers. `rresvport()` is a routine which returns a descriptor to a socket with an address in the privileged port space. `ruserok()` is a routine used by servers to authenticate clients requesting service with `rcmd`. All three functions are present in the same file and are used by the `rshd(8C)` server (among others).

`rcmd()` looks up the host *ahost* using `gethostbyname` (see `gethostent(3N)`), returning `-1` if the host does not exist. Otherwise *ahost* is set to the standard name of the host and a connection is established to a server residing at the well-known Internet port *inport*.

If the connection succeeds, a socket in the Internet domain of type `SOCK_STREAM` is returned to the caller, and given to the remote command as its standard input (file descriptor 0) and standard output (file descriptor 1). If *fd2p* is non-zero, then an auxiliary channel to a control process will be set up, and a descriptor for it will be placed in *fd2p*. The control process will return diagnostic output from the command (file descriptor 2) on this channel, and will also accept bytes on this channel as signal numbers, to be forwarded to the process group of the command. If *fd2p* is 0, then the standard error (file descriptor 2) of the remote command will be made the same as its standard output and no provision is made for sending arbitrary signals to the remote process, although you may be able to get its attention by using out-of-band data.

The protocol is described in detail in `rshd(8C)`.

The `rresvport()` routine is used to obtain a socket with a privileged address bound to it. This socket is suitable for use by `rcmd()` and several other routines. Privileged Internet ports are those in the range 0 to 1023. Only the super-user is allowed to bind an address of this sort to a socket.

`ruserok()` takes a remote host's name, as returned by a `gethostbyaddr` (see `gethostent(3N)`) routine, two user names and a flag indicating whether the local user's name is that of the super-user. It then checks the files `/etc/hosts.equiv` and, possibly, `.rhosts` in the local user's home directory to see if the request for service is allowed. A 0 is returned if the machine name is listed in the `/etc/hosts.equiv` file, or the host and remote user name are found in the `.rhosts` file; otherwise `ruserok()` returns `-1`. If the super-user flag is 1, the checking of the `/etc/hosts.equiv` file is bypassed.

## FILES

```
/etc/hosts.equiv
.rhosts
```

## SEE ALSO

```
rlogin(1C), rsh(1C), intro(2), gethostent(3N), rexec(3N), rexecd(8C), rlogind(8C), rshd(8C)
```

**DIAGNOSTICS**

**rcmd()** returns a valid socket descriptor on success. It returns -1 on error and prints a diagnostic message on the standard error.

**rresvport()** returns a valid, bound socket descriptor on success. It returns -1 on error with the global value **errno** set according to the reason for failure. The error code EAGAIN is overloaded to mean "All network ports in use."



**NAME**

`realpath` – return the canonicalized absolute pathname

**SYNOPSIS**

```
#include <sys/param.h>

char *realpath(path, resolved_path)
char *path;
char resolved_path[MAXPATHLEN];
```

**DESCRIPTION**

`realpath()` expands all symbolic links and resolves references to `'./'`, `'../'` and extra `'/'` characters in the null terminated string named by `path` and stores the canonicalized absolute pathname in the buffer named by `resolved_path`. The resulting path will have no symbolic links components, nor any `'./'` or `'../'` components.

**RETURN VALUES**

`realpath()` returns a pointer to the `resolved_path` on success. On failure, it returns `NULL`, sets `errno` to indicate the error, and places in `resolved_path` the absolute pathname of the `path` component which could not be resolved.

**ERRORS**

<code>EACCES</code>	Search permission is denied for a component of the path prefix of <code>path</code> .
<code>EFAULT</code>	<code>resolved_path</code> extends outside the process's allocated address space.
<code>ELOOP</code>	Too many symbolic links were encountered in translating <code>path</code> .
<code>EINVAL</code>	<code>path</code> or <code>resolved_path</code> was <code>NULL</code> .
<code>EIO</code>	An I/O error occurred while reading from or writing to the file system.
<code>ENAMETOOLONG</code>	The length of the path argument exceeds <code>{PATH_MAX}</code> . A pathname component is longer than <code>{NAME_MAX}</code> (see <code>sysconf(2V)</code> ) while <code>{_POSIX_NO_TRUNC}</code> is in effect (see <code>pathconf(2V)</code> ).
<code>ENOENT</code>	The named file does not exist.

**SEE ALSO**

`readlink(2)`, `getwd(3)`

**WARNINGS**

It indirectly invokes the `readlink(2)` system call and `getwd(3)` library call (for relative path names), and hence inherits the possibility of hanging due to inaccessible file system resources.

**NAME**

regex, re\_comp, re\_exec – regular expression handler

**SYNOPSIS**

**char \*re\_comp(s)**

**char \*s;**

**re\_exec(s)**

**char \*s;**

**DESCRIPTION**

**re\_comp()** compiles a string into an internal form suitable for pattern matching. **re\_exec()** checks the argument string against the last string passed to **re\_comp()**.

**re\_comp()** returns a NULL pointer if the string *s* was compiled successfully; otherwise a string containing an error message is returned. If **re\_comp()** is passed 0 or a null string, it returns without changing the currently compiled regular expression.

**re\_exec()** returns 1 if the string *s* matches the last compiled regular expression, 0 if the string *s* failed to match the last compiled regular expression, and -1 if the compiled regular expression was invalid (indicating an internal error).

The strings passed to both **re\_comp()** and **re\_exec()** may have trailing or embedded NEWLINE characters; they are terminated by null characters. The regular expressions recognized are described in the manual entry for **ed(1)**, given the above difference.

**SEE ALSO**

**ed(1)**, **ex(1)**, **grep(1V)**

**DIAGNOSTICS**

**re\_exec()** returns -1 for an internal error.

**re\_comp()** returns one of the following strings if an error occurs:

**No previous regular expression**

**Regular expression too long**

**unmatched \**

**missing ]**

**too many \(\) pairs**

**unmatched \)**

## NAME

regexp – regular expression compile and match routines

## SYNOPSIS

```
#define INIT <declarations>
#define GETC() <getc code>
#define PEEKC() <peekc code>
#define UNGETC(c) <ungetc code>
#define RETURN(pointer) <return code>
#define ERROR(val) <error code>

#include <regexp.h>

char *compile(instring, expbuf, endbuf, eof)
char *instring, *expbuf, *endbuf;
int eof;

int step(string, expbuf)
char *string, *expbuf;

extern char *loc1, *loc2, *locs;

extern int circf, sed, nbra;
```

## DESCRIPTION

This page describes general-purpose regular expression matching routines.

The interface to this file is unpleasantly complex. Programs that include this file must have the following five macros declared before the '#include <regexp.h>' statement. These macros are used by the *compile* routine.

GETC()	Return the value of the next character in the regular expression pattern. Successive calls to GETC() should return successive characters of the regular expression.
PEEKC()	Return the next character in the regular expression. Successive calls to PEEKC() should return the same character, which should also be the next character returned by GETC().
UNGETC(c)	Returns the argument <i>c</i> by the next call to GETC() or PEEKC(). No more than one character of pushback is ever needed and this character is guaranteed to be the last character read by GETC(). The value of the macro UNGETC(c) is always ignored.
RETURN(pointer)	This macro is used on normal exit of the <i>compile</i> routine. The value of the argument <i>pointer</i> is a pointer to the character after the last character of the compiled regular expression. This is useful to programs that have memory allocation to manage.

## ERRORS

ERROR(val)	This is the abnormal return from the <i>compile()</i> routine. The argument <i>val</i> is an error number (see table below for meanings). This call should never return.
------------	--

ERROR	MEANING
11	Range endpoint too large.
16	Bad number.
25	'\ digit' out of range.
36	Illegal or missing delimiter.
41	No remembered search string.
42	\( \) imbalance.
43	Too many \(.

44	More than 2 numbers given in <code>{ }</code> .
45	<code>}</code> expected after <code>\</code> .
46	First number exceeds second in <code>{ }</code> .
49	<code>[ ]</code> imbalance.
50	Regular expression too long.

The syntax of the `compile()` routine is as follows:

**`compile(instring, expbuf, endbuf, eof)`**

The first parameter *instring* is never used explicitly by the `compile()` routine but is useful for programs that pass down different pointers to input characters. It is sometimes used in the `INIT()` declaration (see below). Programs that call functions to input characters or have characters in an external array can pass down a value of `((char *) 0)` for this parameter.

The next parameter *expbuf* is a character pointer. It points to the place where the compiled regular expression will be placed.

The parameter *endbuf* is one more than the highest address where the compiled regular expression may be placed. If the compiled expression cannot fit in `(endbuf-expbuf)` bytes, a call to `ERROR(50)` is made.

The parameter *eof* is the character that marks the end of the regular expression. For example, in an editor like `ed(1)`, this character would usually be `'/'`.

Each program that includes this file must have a `#define` statement for `INIT()`. This definition will be placed right after the declaration for the function `compile()` and `{` (opening curly brace). It is used for dependent declarations and initializations. Most often it is used to set a register variable to point the beginning of the regular expression so that this register variable can be used in the declarations for `GETC()`, `PEEKC()`, and `UNGETC()`. Otherwise it can be used to declare external variables that might be used by `GETC()`, `PEEKC()`, and `UNGETC()`. See the example below of the declarations taken from `grep(1V)`.

There are other functions in this file that perform actual regular expression matching, one of which is the function `step()`. The call to `step()` is as follows:

**`step(string, expbuf)`**

The first parameter to `step()` is a pointer to a string of characters to be checked for a match. This string should be null-terminated

The second parameter *expbuf* is the compiled regular expression that was obtained by a call of the function `compile`.

The function `step()` returns non-zero if the given string matches the regular expression, and zero if the expressions do not match. If there is a match, two external character pointers are set as a side effect to the call to `step()`. The variable set in `step()` is *loc1*. This is a pointer to the first character that matched the regular expression. The variable *loc2*, which is set by the function `advance()`, points to the character after the last character that matches the regular expression. Thus if the regular expression matches the entire line, *loc1* will point to the first character of *string* and *loc2* will point to the null character at the end of *string*.

`step()` uses the external variable *circf* which is set by `compile()` if the regular expression begins with `^`. If this is set then `step()` will try to match the regular expression to the beginning of the string only. If more than one regular expression is to be compiled before the first is executed the value of *circf* should be saved for each compiled expression and *circf* should be set to that saved value before each call to `step()`.

The function `advance()` is called from `step()` with the same arguments as `step()`. The purpose of `step()` is to step through the *string* argument and call `advance()` until `advance()` returns non-zero indicating a match or until the end of *string* is reached. If one wants to constrain *string* to the beginning of the line in all cases, `step()` need not be called; simply call `advance()`.

When `advance()` encounters a `*` or `\{ \}` sequence in the regular expression, it will advance its pointer to the string to be matched as far as possible and will recursively call itself trying to match the rest of the string to the rest of the regular expression. As long as there is no match, `advance()` will back up along the string until it finds a match or reaches the point in the string that initially matched the `*` or `\{ \}`. It is sometimes desirable to stop this backing up before the initial point in the string is reached. If the external character pointer `locs` is equal to the point in the string at sometime during the backing up process, `advance()` will break out of the loop that backs up and will return zero. This could be used by an editor like `ed(1)` or `sed(1V)` for substitutions done globally (not just the first occurrence, but the whole line) so, for example, expressions like `s/y*/g` do not loop forever.

The additional external variables `sed` and `nbra` are used for special purposes.

#### EXAMPLES

The following is an example of how the regular expression macros and calls could look in a command like `grep(1V)`:

```
#define INIT    register char *sp = instring;
#define GETC() (*sp++)
#define PEEKC()  (*sp)
#define UNGETC(c)  (—sp)
#define RETURN(c)  return;
#define ERROR(c)  regerr()

#include <regexp.h>
...
                                (void) compile(*argv, expbuf, &expbuf[ESIZE], ^0');
...
                                if (step(linebuf, expbuf))
                                    succeed ();
```

#### SEE ALSO

`ed(1)`, `grep(1V)`, `sed(1V)`

#### BUGS

The handling of `circf` is difficult.

## NAME

resolver, res\_mkquery, res\_send, res\_init, dn\_comp, dn\_expand – resolver routines

## SYNOPSIS

```
#include <sys/types.h>
#include <netinet/in.h>
#include <arpa/nameser.h>
#include <resolv.h>

res_mkquery(op, dname, class, type, data, datalen, newrr, buf, buflen)
int op;
char *dname;
int class, type;
char *data;
int datalen;
struct rrec *newrr;
char *buf;
int buflen;

res_send(msg, msglen, answer, anslen)
char *msg;
int msglen;
char *answer;
int anslen;

res_init()

dn_comp(exp_dn, comp_dn, length, dnptrs, lastdnptr)
u_char *exp_dn, *comp_dn;
int length;
u_char **dnptrs, **lastdnptr;

dn_expand(msg, msglen, comp_dn, exp_dn, length)
u_char *msg, *eomorig, *comp_dn, exp_dn;
int length;
```

## DESCRIPTION

These routines are used for making, sending and interpreting packets to Internet domain name servers. You can link a program with the resolver library using the `-lresolv` argument on the linking command line.

Global information that is used by the resolver routines is kept in the variable `_res`. Most of the values have reasonable defaults and can be ignored. Options are a simple bit mask and are OR'ed in to enable. Options stored in `_res.options` are defined in `<resolv.h>` and are as follows.

<code>RES_INIT</code>	True if the initial name server address and default domain name are initialized (that is, <code>res_init()</code> has been called).
<code>RES_DEBUG</code>	Print debugging messages.
<code>RES_AAONLY</code>	Accept authoritative answers only. <code>res_send()</code> continues until it finds an authoritative answer or finds an error. Currently this is not implemented.
<code>RES_USEVC</code>	Use TCP connections for queries instead of UDP.
<code>RES_STAYOPEN</code>	Used with <code>RES_USEVC</code> to keep the TCP connection open between queries. This is useful only in programs that regularly do many queries. UDP should be the normal mode used.
<code>RES_IGNTC</code>	Unused currently (ignore truncation errors, that is, do not retry with TCP).
<code>RES_RECURSE</code>	Set the recursion desired bit in queries. This is the default. <code>res_send()</code> does not do iterative queries and expects the name server to handle recursion.

**RES\_DEFNAMES** Append the default domain name to single label queries. This is the default.  
**RES\_DNSRCH** Search up the domain tree from the default domain, in all but the top level. This is the default.

**res\_init()** reads the initialization file to get the default domain name and the Internet addresses of the initial name servers. If no **nameserver** line exists, the host running the resolver is tried. **res\_mkquery()** makes a standard query message and places it in *buf*. **res\_mkquery()** returns the size of the query or **-1** if the query is larger than *buflen*. *op* is usually **QUERY** but can be any of the query types defined in **<nameser.h>**. *dname* is the domain name. If *dname* consists of a single label and the **RES\_DEFNAMES** flag is enabled (the default), *dname* is appended with the current domain name. The current domain name is defined in a system file and can be overridden by the environment variable **LOCALDOMAIN**. *newrr* is currently unused but is intended for making update messages.

**res\_send()** sends a query to name servers and returns an answer. It calls **res\_init()** if **RES\_INIT** is not set, send the query to the local name server, and handle timeouts and retries. The length of the message is returned or **-1** if there were errors.

**dn\_expand()** Expands the compressed domain name *comp\_dn* to a full domain name. Expanded names are converted to upper case. *msg* is a pointer to the beginning of the message, *exp\_dn* is a pointer to a buffer of size *length* for the result. The size of compressed name is returned or **-1** if there was an error.

**dn\_comp()** Compresses the domain name *exp\_dn* and stores it in *comp\_dn*. The size of the compressed name is returned or **-1** if there were errors. *length* is the size of the array pointed to by *comp\_dn*. *dnptrs* is a list of pointers to previously compressed names in the current message. The first pointer points to the beginning of the message and the list ends with **NULL**. *lastdnptr* is a pointer to the end of the array pointed to *dnptrs*. A side effect is to update the list of pointers for labels inserted into the message by **dn\_comp()** as the name is compressed. If *dnptr* is **NULL**, do not try to compress names. If *lastdnptr* is **NULL**, do not update the list.

#### FILES

**/etc/resolv.conf** see **resolv.conf(5)**  
**/usr/lib/libresolv.a**

#### SEE ALSO

**resolv.conf(5)**, **named(8C)**  
*System and Network Administration*

#### NOTES

**/usr/lib/libresolv.a** is necessary for compiling programs.

**NAME**

rexec – return stream to a remote command

**SYNOPSIS**

```
rem = rexec(ahost, inport, user, passwd, cmd, fd2p);
char **ahost;
u_short inport;
char *user, *passwd, *cmd;
int *fd2p;
```

**DESCRIPTION**

rexec() looks up the host *\*ahost* using `gethostbyname()` (see `gethostent(3N)`), returning `-1` if the host does not exist. Otherwise *\*ahost* is set to the standard name of the host. If a username and password are both specified, then these are used to authenticate to the foreign host; otherwise the environment and then the user's `.netrc` file in his home directory are searched for appropriate information. If all this fails, the user is prompted for the information.

The port `inport` specifies which well-known DARPA Internet port to use for the connection; it will normally be the value returned from the call `'getservbyname("exec", "tcp")'` (see `getservent(3N)`). The protocol for connection is described in detail in `rexecd(8C)`.

If the call succeeds, a socket of type `SOCK_STREAM` is returned to the caller, and given to the remote command as its standard input and standard output. If *fd2p* is non-zero, then a auxiliary channel to a control process will be setup, and a descriptor for it will be placed in *\*fd2p*. The control process will return diagnostic output from the command (unit 2) on this channel, and will also accept bytes on this channel as signal numbers, to be forwarded to the process group of the command. If *fd2p* is 0, then the standard error (unit 2 of the remote command) will be made the same as its standard output and no provision is made for sending arbitrary signals to the remote process, although you may be able to get its attention by using out-of-band data.

**SEE ALSO**

`gethostent(3N)`, `getservent(3N)`, `rcmd(3N)`, `rexecd(8C)`

**BUGS**

There is no way to specify options to the `socket()` call that `rexec()` makes.



**NAME**

**rpc** – library routines for remote procedure calls

**SYNOPSIS AND DESCRIPTION**

RPC routines allow C programs to make procedure calls on other machines across the network. First, the client calls a procedure to send a request to the server. Upon receipt of the request, the server calls a dispatch routine to perform the requested service, and then sends back a reply. Finally, the procedure call returns to the client.

All RPC routines require the header `<rpc/rpc.h>` to be included.

The RPC routines have been grouped by usage on the following man pages.

- portmap(3N)** Library routines for the RPC bind service, **portmap(8C)**. The routines documented on this page include:
- `pmap_getmaps()`
  - `pmap_getport()`
  - `pmap_rmtcall()`
  - `pmap_set()`
  - `pmap_unset()`
  - `xdr_pmap()`
  - `xdr_pmaplist()`
- rpc\_clnt\_auth(3N)** Library routines for client side remote procedure call authentication. The routines documented on this page include:
- `auth_destroy()`
  - `authnone_create()`
  - `authunix_create()`
  - `authunix_create_default()`
- rpc\_clnt\_calls(3N)** Library routines for client side calls. The routines documented on this page include:
- `callrpc()`
  - `clnt_broadcast()`
  - `clnt_call()`
  - `clnt_freeres()`
  - `clnt_geterr()`
  - `clnt_perrno()`
  - `clnt_perror()`
  - `clnt_sperrno()`
  - `clnt_sperror()`
- rpc\_clnt\_create(3N)** Library routines for dealing with the creation and manipulation of CLIENT handles. The routines documented on this page include:
- `clnt_control()`
  - `clnt_create()`
  - `clnt_create_vers()`
  - `clnt_destroy()`
  - `clnt_pcreateerror()`
  - `clntraw_create()`
  - `clnt_spcreateerror()`
  - `clnttcp_create()`
  - `clntudp_bufcreate()`
  - `clntudp_create()`
  - `rpc_createrr()`

- rpc\_svc\_calls(3N)** Library routines for registering servers. The routines documented on this page include:
- registerrpc()**
  - svc\_register()**
  - svc\_unregister()**
  - xprt\_register()**
  - xprt\_unregister()**
- rpc\_svc\_create(3N)** Library routines for dealing with the creation of server side handles. The routines documented on this page include:
- svc\_destroy()**
  - svcfld\_create()**
  - svcrow\_create()**
  - svctcp\_create()**
  - svcupd\_bufcreate()**
- rpc\_svc\_err(3N)** Library routines for server side remote procedure call errors. The routines documented on this page include:
- svcerr\_auth()**
  - svcerr\_decode()**
  - svcerr\_noproc()**
  - svcerr\_noprogram()**
  - svcerr\_progvers()**
  - svcerr\_systemerr()**
  - svcerr\_weakauth()**
- rpc\_svc\_reg(3N)** Library routines for RPC servers. The routines documented on this page include:
- svc\_fds()**
  - svc\_fdset()**
  - svc\_freeargs()**
  - svc\_getargs()**
  - svc\_getcaller()**
  - svc\_getreq()**
  - svc\_getreqset()**
  - svc\_run()**
  - svc\_sendreply()**
- rpc\_xdr(3N)** XDR library routines for remote procedure calls. The routines documented on this page include:
- xdr\_accepted\_reply()**
  - xdr\_authunix\_parms()**
  - xdr\_callhdr()**
  - xdr\_callmsg()**
  - xdr\_opaque\_auth()**
  - xdr\_rejected\_reply()**
  - xdr\_replymsg()**

**secure\_rpc(3N)**

Library routines for secure remote procedure calls. The routines documented on this page include:

- authdes\_create()**
- authdes\_getucred()**
- get\_mayaddress()**
- getnetname()**
- host2netname()**
- key\_decryptsession()**
- key\_encryptsession()**
- key\_gendes()**
- key\_setsecret()**
- netname2host()**
- netname2user()**
- user2netname()**

**SEE ALSO**

**portmap(3N), rpc\_clnt\_auth(3N), rpc\_clnt\_calls(3N), rpc\_clnt\_create(3N), rpc\_svc\_calls(3N), rpc\_svc\_create(3N), rpc\_svc\_err(3N), rpc\_svc\_reg(3N), rpc\_xdr(3N), secure\_rpc(3N), xdr(3N), publickey(5), portmap(8C), keyser(8C)**

*Network Programming*

**NAME**

`auth_destroy`, `authnone_create`, `authunix_create`, `authunix_create_default` – library routines for client side remote procedure call authentication

**DESCRIPTION**

RPC routines allow C programs to make procedure calls on other machines across the network. First, the client calls a procedure to send a request to the server. Upon receipt of the request, the server calls a dispatch routine to perform the requested service, and then sends back a reply. Finally, the procedure call returns to the client.

RPC allows various authentication types. Currently, it supports `AUTH_NONE`, `AUTH_UNIX`, `AUTH_DES`. For routines relating to the `AUTH_DES` type, see `secure_rpc(3N)`.

These routines are called after creating the `CLIENT` handle. The client's authentication information is passed to the server when the RPC call is made.

**Routines**

The following routines require that the header `<rpc.h>` be included. The `AUTH` data structure is defined in the RPC/XDR Library Definitions of the *Network Programming*.

```
#include <rpc/rpc.h>
```

```
void auth_destroy(auth)
AUTH *auth;
```

Destroy the authentication information associated with *auth*. Destruction usually involves deallocation of private data structures. The use of *auth* is undefined after calling `auth_destroy()`.

```
AUTH * authnone_create()
```

Create and return an RPC authentication handle that passes no usable authentication information with each remote procedure call. This is the default authentication used by RPC.

```
AUTH * authunix_create(host, uid, gid, grouplen, gidlistp)
char *host;
int uid, gid, grouplen, *gidlistp;
```

Create and return an RPC authentication handle that contains authentication information. The parameter *host* is the name of the machine on which the information was created; *uid* is the user's user ID; *gid* is the user's current group ID; *grouplen* and *gidlistp* refer to a counted array of groups to which the user belongs. Warning: It is not very difficult to impersonate a user.

```
AUTH * authunix_create_default()
```

Call `authunix_create()` with the appropriate parameters.

**SEE ALSO**

`rpc(3N)`, `rpc_clnt_create(3N)`, `rpc_clnt_calls(3N)`

**NAME**

callrpc, clnt\_broadcast, clnt\_call, clnt\_freeres, clnt\_geterr, clnt\_perrno, clnt\_perror, clnt\_sperno, clnt\_spperror – library routines for client side calls

**DESCRIPTION**

RPC routines allow C programs to make procedure calls on other machines across the network. First, the client calls a procedure to send a request to the server. Upon receipt of the request, the server calls a dispatch routine to perform the requested service, and then sends back a reply. Finally, the procedure call returns to the client.

The `clnt_call()`, `callrpc()` and `clnt_broadcast()` routines handle the client side of the procedure call. The remaining routines deal with error handling in the case of errors.

**Routines**

The CLIENT data structure is defined in the RPC/XDR Library Definition of the *Network Programming*.

```
#include <rpc/rpc.h>
```

```
int callrpc(host, prognum, versnum, procnum, inproc, in, outproc, out)
char *host;
u_long prognum, versnum, procnum;
char *in;
xdrproc_t inproc;
char *out;
xdrproc_t outproc;
```

Call the remote procedure associated with *prognum*, *versnum*, and *procnum* on the machine, *host*. The parameter *in* is the address of the procedure's argument, and *out* is the address of where to place the result; *inproc* is an XDR function used to encode the procedure's parameters, and *outproc* is an XDR function used to decode the procedure's results. This routine returns 0 if it succeeds, or the value of enum `clnt_stat` cast to an integer if it fails. Use `clnt_perrno()` to translate failure statuses into messages.

Warning: Calling remote procedures with this routine uses UDP/IP as the transport; see `clntudp_create()` on `rpc_clnt_create(3N)` for restrictions. You do not have control of timeouts or authentication using this routine.

```
enum clnt_stat clnt_broadcast(prognum, versnum, procnum, inproc, in, outproc, out, eachresult)
u_long prognum, versnum, procnum;
char *in;
xdrproc_t inproc;
char *out;
xdrproc_t outproc;
bool_t eachresult;
```

Like `callrpc()`, except the call message is broadcast to all locally connected broadcast nets. Each time the caller receives a response, this routine calls `eachresult()`, whose form is:

```
int eachresult(out, addr)
char *out;
struct sockaddr_in *addr;
```

where *out* is the same as *out* passed to `clnt_broadcast()`, except that the remote procedure's output is decoded there; *addr* points to the address of the machine that sent the results. If `eachresult()` returns 0 `clnt_broadcast()` waits for more replies; otherwise it returns with appropriate status. If `eachresult()` is NULL, `clnt_broadcast()` returns without waiting for any replies.

Note: `clnt_broadcast()` uses AUTH\_UNIX style of authentication.

Warning: Broadcast packets are limited in size to the maximum transfer unit of the data link. For Ethernet, the callers argument size should not exceed 1400 bytes.

```
enum clnt_stat clnt_call(clnt, procnum, inproc, in, outproc, out, timeout)
CLIENT *clnt;
u_long procnum;
xdrproc_t inproc, outproc;
char *in, *out;
struct timeval timeout;
```

Call the remote procedure *procnum* associated with the client handle, *clnt*, which is obtained with an RPC client creation routine such as `clnt_create()` (see `rpc_clnt_create(3N)`). The parameter *in* is the address of the procedure's argument, and *out* is the address of where to place the result; *inproc* is an XDR function used to encode the procedure's parameters in XDR, and *outproc* is used to decode the procedure's results; *timeout* is the time allowed for a response from the server.

```
bool_t clnt_freeres(clnt, outproc, out)
CLIENT *clnt;
xdrproc_t outproc;
char *out;
```

Free any data allocated by the RPC/XDR system when it decoded the results of an RPC call. The parameter *out* is the address of the results, and *outproc* is the XDR routine describing the results. This routine returns TRUE if the results were successfully freed, and FALSE otherwise. Note: This is equivalent to doing `xdr_free(outproc, out)` (see `xdr_simple(3N)`).

```
void clnt_geterr(clnt, errp)
CLIENT *clnt;
struct rpc_err *errp;
```

Copy the error structure out of the client handle to the structure at address *errp*. *errp* should point to preallocated space.

```
void clnt_perrno(stat)
enum clnt_stat stat;
```

Print a message to the standard error corresponding to the condition indicated by *stat*. A NEWLINE is appended at the end of the message. Used after `callrpc()` or `clnt_broadcast()`.

```
void clnt_perror(clnt, str)
CLIENT *clnt;
char *str;
```

Print a message to the standard error indicating why an RPC call failed; *clnt* is the handle used to do the call. The message is prepended with string *s* and a colon. A NEWLINE is appended at the end of the message. Used after `clnt_call()`.

```
char *clnt_sperrno(stat)
enum clnt_stat stat;
```

Take the same arguments as `clnt_perrno()`, but instead of sending a message to the standard error indicating why an RPC failed, return a pointer to a string which contains the message. `clnt_sperrno()` does not append a NEWLINE at the end of the message.

`clnt_sperrno()` is used instead of `clnt_perrno()` if the program does not have a standard error (as a program running as a server quite likely does not), or if the programmer does not want the message to be output with `printf(3V)`, or if a message format different than that supported by `clnt_perrno()` is to be used.

```
char *clnt_sperror(clnt, str)
CLIENT *clnt;
char *str;
```

Like `clnt_perror()`, except that (like `clnt_sperrno()`) it returns a string instead of printing to the standard error. Unlike `clnt_perror()`, it does not append the message with a NEWLINE.

Note: `clnt_sperror()` returns pointer to a static buffer that is overwritten on each call.

**SEE ALSO**

`printf(3V)`, `rpc(3N)`, `rpc_clnt_auth(3N)`, `rpc_clnt_create(3N)`, `xdr_simple(3N)`

## NAME

clnt\_control, clnt\_create, clnt\_create\_vers, clnt\_destroy, clnt\_pcreateerror, clntraw\_create, clnt\_spcreateerror, clnttcp\_create, clntudp\_bufcreate, rpc\_createrr – library routines for dealing with creation and manipulation of CLIENT handles

## DESCRIPTION

RPC routines allow C programs to make procedure calls on other machines across the network. First, the client calls a procedure to send a request to the server. Upon receipt of the request, the server calls a dispatch routine to perform the requested service, and then sends back a reply. Finally, the procedure call returns to the client.

The CLIENT data structure is defined in the RPC/XDR Library Definition of the *Network Programming*.

```
#include <rpc/rpc.h>
```

```
bool_t clnt_control(clnt, request, info)
```

```
CLIENT *clnt;
```

```
int request;
```

```
char *info;
```

Change or retrieve various information about a client object. *request* indicates the type of operation, and *info* is a pointer to the information. For both UDP and TCP, the supported values of *request* and their argument types and what they do are:

CLSET_TIMEOUT	struct timeval	set total timeout
CLGET_TIMEOUT	struct timeval	get total timeout
CLGET_FD	int	get associated socket
CLSET_FD_CLOSE	void	close socket on <code>clnt_destroy()</code>
CLSET_FD_NCLOSE	void	leave socket open on <code>clnt_destroy()</code>

Note: If you set the timeout using `clnt_control()`, the timeout parameter passed to `clnt_call()` (see `rpc_clnt_calls(3N)`) will be ignored in all future calls.

CLGET_SERVER_ADDR	struct sockaddr_in	get server's address
-------------------	--------------------	----------------------

The following operations are valid for UDP only:

CLSET_RETRY_TIMEOUT	struct timeval	set the retry timeout
CLGET_RETRY_TIMEOUT	struct timeval	get the retry timeout

The retry timeout is the time that UDP RPC waits for the server to reply before retransmitting the request.

This routine returns TRUE on success, and FALSE on failure.

```
CLIENT * clnt_create(host, prognum, versnum, protocol)
```

```
char *host;
```

```
u_long prognum, versnum;
```

```
char *protocol;
```

Generic client creation routine for program *prognum* and version *versnum*. *host* identifies the name of the remote host where the server is located. *protocol* indicates which kind of transport protocol to use. The currently supported values for this field are "udp" and "tcp". Default timeouts are set, but they can be modified using `clnt_control()`. If successful it returns a client handle, otherwise it returns NULL.



Warning: Using UDP has its shortcomings. Since UDP-based RPC messages can only hold up to 8 Kbytes of encoded data, this transport cannot be used for procedures that take arguments or return results larger than 8 Kbytes. Use TCP instead.

Note: If the requested version number *versnum* is not registered with the `portmap(8C)` service on *host*, but at least a version number for the given program number is registered, `clnt_create()` returns a handle. The version mismatch will be discovered by a `clnt_call()` later (see `rpc_clnt_calls(3N)`).

```
CLIENT * clnt_create_vers(host, prognum, vers_outp, vers_low, vers_high, protocol)
char *host;
u_long prognum;
u_long *vers_outp;
u_long vers_low, vers_high;
char *protocol;
```

This is a generic client creation routine which also checks for the version available. *host* identifies the name of the remote host where the server is located. *protocol* indicates which kind of transport protocol to use. The currently supported values for this field are "udp" and "tcp". If the routine is successful it returns a client handle created for the highest version between *vers\_low* and *vers\_high* that is supported by the server. *vers\_outp* is set to this value. That is, after a successful return  $vers\_low \leq *vers\_outp \leq vers\_high$ . If no version between *vers\_low* and *vers\_high* is supported by the server then the routine fails and returns NULL. Default timeouts are set, but can be modified using `clnt_control()`.

Note: `clnt_create()` returns a valid client handle even if the particular version number supplied to `clnt_create()` is not registered with the portmap service. This mismatch will be discovered by a `clnt_call()` later (see `rpc_clnt_calls(3N)`). However, `clnt_create_vers()` does this for you and returns a valid handle only if a version within the range supplied is supported by the server.

```
void clnt_destroy(clnt)
CLIENT *clnt;
```

Destroy the client's RPC handle. Destruction usually involves deallocation of private data structures, including *clnt* itself. Use of *clnt* is undefined after calling `clnt_destroy()`. If the RPC library opened the associated socket, or `CLSET_FD_CLOSE` was set using `clnt_control()`, `clnt_destroy()` closes the socket.

```
void clnt_pcreateerror(str)
char *str;
```

Print a message to the standard error indicating why a client handle could not be created. The message is prepended with string *s* and a colon. Used when routines such as `clnt_create()`, `clntraw_create()`, `clnttcp_create()`, or `clntudp_create()` fails.

```
CLIENT * clntraw_create(prognum, versnum)
u_long prognum, versnum;
```

Create an RPC client for the remote program *prognum*, version *versnum*. The transport used to pass messages to the service is actually a buffer within the process's address space, so the corresponding RPC server should live in the same address space; also see `svcrw_create()` (see `rpc_svc_create(3N)`). This allows simulation of RPC and getting RPC overheads, such as round trip times, without any kernel interference. If successful it returns a client handle, otherwise it returns NULL.

```
char * clnt_screateerror(str)
char *str;
```

Like `clnt_pcreateerror()`, except that it returns a string instead of printing to the standard error. It, however, does not append the message with a `NEWLINE`.

Note: `clnt_screateerror()` returns a pointer to a static buffer that is overwritten on each call.

```
CLIENT * clnttcp_create(addr, prognum, versnum, sockp, sendsz, recvsz)
struct sockaddr_in *addr;
u_long prognum, versnum;
int *sockp;
u_int sendsz, recvsz;
```

Create a client handle for the remote program *prognum*, version *versnum*; the client uses TCP/IP as a transport. The remote program is located at Internet address *addr*. If `addr->sin_port` is zero, it is set to the port on which the remote program is listening (the remote `portmap` service is consulted for this information). The parameter *sockp* is a pointer to a socket; if it is `RPC_ANYSOCK`, then a new socket is opened and *sockp* is updated. Since TCP-based RPC uses buffered I/O, the user may specify the size of the send and receive buffers with the parameters *sendsz* and *recvsz*; values of zero choose defaults. If successful it returns a client handle, otherwise it returns `NULL`.

Warning: If `addr->sin_port` is zero and the requested version number *versnum* is not registered with the remote `portmap` service, it returns a handle if at least a version number for the given program number is registered. The version mismatch will be discovered by a `clnt_call()` later (see `rpc_clnt_calls(3N)`).

```
CLIENT * clntudp_bufcreate(addr, prognum, versnum, wait, sockp, sendsz, recvsz)
struct sockaddr_in *addr;
u_long prognum, versnum;
struct timeval wait;
int *sockp;
u_int sendsz;
u_int recvsz;
```

Create a client handle for the remote program *prognum*, on *versnum*; the client uses UDP/IP as the transport. The remote program is located at the Internet address *addr*. If `addr->sin_port` is zero, it is set to port on which the remote program is listening on (the remote `portmap` service is consulted for this information). The parameter *sockp* is a pointer to a socket; if it is `RPC_ANYSOCK`, then a new socket is opened and *sockp* is updated. The UDP transport resends the call message in intervals of *wait* time until a response is received or until the call times out. The total time for the call to time out is specified by `clnt_call()` (see `rpc_clnt_calls(3N)`). If successful it returns a client handle, otherwise it returns `NULL`.

The user can specify the maximum packet size for sending and receiving by using *sendsz* and *recvsz* arguments for UDP-based RPC messages.

Warning: If `addr->sin_port` is zero and the requested version number *versnum* is not registered with the remote `portmap` service, it returns a handle if at least a version number for the given program number is registered. The version mismatch is discovered by a `clnt_call()` later (see `rpc_clnt_calls(3N)`).

```
CLIENT * clntudp_create(addr, prognum, versnum, wait, sockp)
struct sockaddr_in *addr;
u_long prognum, versnum;
struct timeval wait;
int *sockp;
```

Create a client handle for the remote program *prognum*, version *versnum*; the client uses UDP/IP as the transport. The remote program is located at the Internet address *addr*. If *addr->sin\_port* is zero, then it is set to actual port that the remote program is listening on (the remote **portmap** service is consulted for this information). The parameter *sockp* is a pointer to a socket; if it is **RPC\_ANYSOCK**, a new socket is opened and *sockp* is updated. The UDP transport resends the call message in intervals of *wait* time until a response is received or until the call times out. The total time for the call to time out is specified by **clnt\_call()** (see **rpc\_clnt\_calls(3N)**). If successful it returns a client handle, otherwise it returns NULL.

Warning: Since UDP-based RPC messages can only hold up to 8 Kbytes of encoded data, this transport cannot be used for procedures that take arguments or results larger than 8 Kbytes. TCP should be used instead.

Warning: If *addr->sin\_port* is zero and the requested version number *versnum* is not registered with the remote portmap service, it returns a handle if any version number for the given program number is registered. The version mismatch is discovered by a **clnt\_call()** later (see **rpc\_clnt\_calls(3N)**).

```
struct rpc_createerr rpc_createerr;
```

A global variable whose value is set by any RPC client handle creation routine that fails. It is used by the routine **clnt\_pcreateerror()** to print the reason for the failure.

**SEE ALSO**

**portmap(3N), rpc(3N), rpc\_clnt\_auth(3N), rpc\_clnt\_calls(3N), rpc\_svc\_create(3N)**

## NAME

registerrpc, svc\_register, svc\_unregister, xpvt\_register, xpvt\_unregister – library routines for registering servers

## DESCRIPTION

These routines are a part of the RPC library which allows the RPC servers to register themselves with portmap(8C), and it associates the given program and version number with the dispatch function.

## Routines

The SVCXPRT data structure is defined in the RPC/XDR Library Definition of the *Network Programming*.

```
#include <rpc/rpc.h>
```

```
int registerrpc(prognum, versnum, procnum, procname, inproc, outproc)
```

```
u_long prognum, versnum, procnum;
```

```
char>(*procname) ();
```

```
xdrproc_t inproc, outproc;
```

Register procedure *procname* with the RPC service package. If a request arrives for program *prognum*, version *versnum*, and procedure *procnum*, *procname* is called with a pointer to its parameter; *procname* must be a procedure that returns a pointer to its static result; *inproc* is used to decode the parameters while *outproc* is used to encode the results. This routine returns 0 if the registration succeeded, -1 otherwise.

Warning: Remote procedures registered in this form are accessed using the UDP/IP transport; see `svcudp_create()` on `rpc_svc_create(3N)` for restrictions. This routine should not be used more than once for the same program and version number.

```
bool_t svc_register(xprt, prognum, versnum, dispatch, protocol)
```

```
SVCXPRT *xprt;
```

```
u_long prognum, versnum;
```

```
void (*dispatch) ();
```

```
u_long protocol;
```

Associates *prognum* and *versnum* with the service dispatch procedure, *dispatch*. If *protocol* is zero, the service is not registered with the portmap service. If *protocol* is non-zero, a mapping of the triple [*prognum*, *versnum*, *protocol*] to *xprt*→*xp\_port* is established with the local portmap service (generally *protocol* is zero, IPPROTO\_UDP or IPPROTO\_TCP). The procedure *dispatch* has the following form:

```
dispatch(request, xprt)
struct svc_req *request;
SVCXPRT *xprt;
```

The `svc_register()` routine returns TRUE if it succeeds, and FALSE otherwise.

```
void svc_unregister(prognum, versnum)
```

```
u_long prognum, versnum;
```

Remove all mapping of the pair [*prognum*, *versnum*] to dispatch routines, and of the triple [*prognum*, *versnum*, \*] to port number.

```
void xpvt_register(xprt)
```

```
SVCXPRT *xprt;
```

After RPC service transport handles are created, they should register themselves with the RPC service package. This routine modifies the global variable `svc_fds`. Service implementors usually do not need this routine.

```
void xpvt_unregister(xprt)  
SVCXPRT *xpvt;
```

Before an RPC service transport handle is destroyed, it should unregister itself with the RPC service package. This routine modifies the global variable `svc_fds`. Service implementors usually do not need this routine directly.

**SEE ALSO**

**portmap(3N), rpc(3N), rpc\_svc\_err(3N), rpc\_svc\_create(3N), rpc\_svc\_reg(3N), portmap(8C)**

**NAME**

`svc_destroy`, `svcfld_create`, `svccraw_create`, `svctcp_create`, `svcudp_bufcreate` – library routines for dealing with the creation of server handles

**DESCRIPTION**

RPC routines allow C programs to make procedure calls on other machines across the network. First, the client calls a procedure to send a request to the server. Upon receipt of the request, the server calls a dispatch routine to perform the requested service, and then sends back a reply. Finally, the procedure call returns to the client.

The `SVCXPRT` data structure is defined in the RPC/XDR Library Definitions of the *Network Programming*.

```
#include <rpc/rpc.h>
```

```
void svc_destroy(xprt)
SVCXPRT *xprt;
```

Destroy the RPC service transport handle, `xprt`. Destruction usually involves deallocation of private data structures, including `xprt` itself. Use of `xprt` is undefined after calling this routine.

```
SVCXPRT * svcfld_create(fd, sendsz, recvsz)
int fd;
u_int sendsz;
u_int recvsz;
```

Create a service on top of any open and bound descriptor and return the handle to it. Typically, this descriptor is a connected socket for a stream protocol such as TCP. `sendsz` and `recvsz` indicate sizes for the send and receive buffers. If they are zero, a reasonable default is chosen. It returns NULL if it fails.

```
SVCXPRT * svccraw_create()
```

This routine creates a RPC service transport, to which it returns a pointer. The transport is a buffer within the process's address space, so the corresponding RPC client must live in the same address space; see `clntraw_create()` on `rpc_clnt_create(3N)`. This routine allows simulation of RPC and getting RPC overheads (such as round trip times), without any kernel interference. This routine returns NULL if it fails.

```
SVCXPRT * svctcp_create(sock, sendsz, recvsz)
int sock;
u_int sendsz, recvsz;
```

This routine creates a TCP/IP-based RPC service transport, to which it returns a pointer. The transport is associated with the socket `sock`. If `sock` is `RPC_ANYSOCK`, then a new socket is created. If the socket is not bound to a local TCP port, then this routine binds it to an arbitrary port. Upon completion, `xprt->xp_sock` is the transport's socket descriptor, and `xprt->xp_port` is the port number on which it is listening. This routine returns NULL if it fails. Since TCP-based RPC uses buffered I/O, users may specify the size of buffers with `sendsz` and `recvsz`; values of zero choose defaults.

```
SVCXPRT * svcudp_bufcreate(sock, sendsz, recvsz)
int sock;
u_int sendsz, recvsz;
```

This routine creates a UDP/IP-based RPC service transport, to which it returns a pointer. The transport is associated with the socket *sock*. If *sock* is `RPC_ANYSOCK`, then a new socket is created. If the socket is not bound to a local UDP port, then this routine binds it to an arbitrary port. Upon completion, `xprt->xp_sock` is the service's socket descriptor, and `xprt->xp_port` is the service's port number. This routine returns NULL if it fails.

The user specifies the maximum packet size for sending and receiving UDP-based RPC messages by using the *sendsz* and *recvsz* parameters.

**SEE ALSO**

`rpc(3N)`, `rpc_clnt_create(3N)`, `rpc_svc_calls(3N)`, `rpc_svc_err(3N)`, `rpc_svc_reg(3N)`, `portmap(8C)`

**NAME**

svcerr\_auth, svcerr\_decode, svcerr\_noproc, svcerr\_noprog, svcerr\_progvers, svcerr\_systemerr, svcerr\_weakauth – library routines for server side remote procedure call errors

**DESCRIPTION**

RPC routines allow C programs to make procedure calls on other machines across the network. First, the client calls a procedure to send a request to the server. Upon receipt of the request, the server calls a dispatch routine to perform the requested service, and then sends back a reply. Finally, the procedure call returns to the client.

These routines can be called by the server side dispatch function if there is any error in the transaction with the client.

**Routines**

The SVCXPRT data structure is defined in the RPC/XDR Library Definitions of the *Network Programming*.

```
#include <rpc/rpc.h>
```

```
void svcerr_auth(xprt, why)
SVCXPRT *xprt;
enum auth_stat why;
```

Called by a service dispatch routine that refuses to perform a remote procedure call due to an authentication error.

```
void svcerr_decode(xprt)
SVCXPRT *xprt;
```

Called by a service dispatch routine that cannot successfully decode the remote parameters. See `svc_getargs()` in `rpc_svc_reg(3N)`.

```
void svcerr_noproc(xprt)
SVCXPRT *xprt;
```

Called by a service dispatch routine that does not implement the procedure number that the caller requests.

```
void svcerr_noprog(xprt)
SVCXPRT *xprt;
```

Called when the desired program is not registered with the RPC package. Service implementors usually do not need this routine.

```
void svcerr_progvers(xprt)
SVCXPRT *xprt;
```

Called when the desired version of a program is not registered with the RPC package. Service implementors usually do not need this routine.

```
void svcerr_systemerr(xprt)
SVCXPRT *xprt;
```

Called by a service dispatch routine when it detects a system error not covered by any particular protocol. For example, if a service can no longer allocate storage, it may call this routine.



```
void svcerr_weakauth(xprt)
SVCXPRT *xprt;
```

Called by a service dispatch routine that refuses to perform a remote procedure call due to insufficient authentication parameters. The routine calls `svcerr_auth(xprt, AUTH_TOOWEAK)`.

**SEE ALSO**

`rpc(3N)`, `rpc_svc_calls(3N)`, `rpc_svc_create(3N)`, `rpc_svc_reg(3N)`

**NAME**

svc\_fds, svc\_fdset, svc\_freeargs, svc\_getargs, svc\_getcaller, svc\_getreq, svc\_getreqset, svc\_getcaller, svc\_run, svc\_sendreply – library routines for RPC servers

**DESCRIPTION**

RPC routines allow C programs to make procedure calls on other machines across the network. First, the client calls a procedure to send a request to the server. Upon receipt of the request, the server calls a dispatch routine to perform the requested service, and then sends back a reply. Finally, the procedure call returns to the client.

These routines are associated with the server side of the RPC mechanism. Some of them are called by the server side dispatch function, while others (such as `svc_run()`) are called when the server is initiated.

**Routines**

The `SVCXPRT` data structure is defined in the RPC/XDR Library Definitions of the *Network Programming*.

```
#include <rpc/rpc.h>
```

```
int svc_fds;
```

Similar to `svc_fdset`, but limited to 32 descriptors. This interface is obsoleted by `svc_fdset`.

```
fd_set svc_fdset;
```

A global variable reflecting the RPC server's read file descriptor bit mask; it is suitable as a parameter to the `select()` system call. This is only of interest if a service implementor does not call `svc_run()`, but rather does their own asynchronous event processing. This variable is read-only (do not pass its address to `select()`!), yet it may change after calls to `svc_getreqset()` or any creation routines.

```
bool_t svc_freeargs(xprt, inproc, in)
```

```
SVCXPRT *xprt;
```

```
xdrproc_t inproc;
```

```
char *in;
```

Free any data allocated by the RPC/XDR system when it decoded the arguments to a service procedure using `svc_getargs()`. This routine returns TRUE if the results were successfully freed, and FALSE otherwise.

```
bool_t svc_getargs(xprt, inproc, in)
```

```
SVCXPRT *xprt;
```

```
xdrproc_t inproc;
```

```
char *in;
```

Decode the arguments of an RPC request associated with the RPC service transport handle, `xprt`. The parameter `in` is the address where the arguments will be placed; `inproc` is the XDR routine used to decode the arguments. This routine returns TRUE if decoding succeeds, and FALSE otherwise.

```
struct sockaddr_in * svc_getcaller(xprt)
```

```
SVCXPRT *xprt;
```

The approved way of getting the network address of the caller of a procedure associated with the RPC service transport handle, `xprt`.

**void svc\_getreq(rdfds)**  
**int rdfds;**

Similar to `svc_getreqset()`, but limited to 32 descriptors. This interface is obsolete by `svc_getreqset()`.

**void svc\_getreqset(rdfdsp)**  
**fd\_set \*rdfdsp;**

This routine is only of interest if a service implementor does not use `svc_run()`, but instead implements custom asynchronous event processing. It is called when the `select()` system call has determined that an RPC request has arrived on some RPC socket(s); `rdfdsp` is the resultant read file descriptor bit mask. The routine returns when all sockets associated with the value of `rdfdsp` have been serviced.

**void svc\_run()**

Normally, this routine only returns in the case of some errors. It waits for RPC requests to arrive, and calls the appropriate service procedure using `svc_getreq()` when one arrives. This procedure is usually waiting for a `select()` system call to return.

**bool\_t svc\_sendreply(xprt, outproc, out)**  
**SVCXPRT \*xprt;**  
**xdrproc\_t outproc;**  
**char \*out;**

Called by an RPC service's dispatch routine to send the results of a remote procedure call. The parameter `xprt` is the request's associated transport handle; `outproc` is the XDR routine which is used to encode the results; and `out` is the address of the results. This routine returns TRUE if it succeeds, FALSE otherwise.

**SEE ALSO**

`select(2)`, `rpc(3N)`, `rpc_svc_calls(3N)`, `rpc_svc_create(3N)`, `rpc_svc_err(3N)`

**NAME**

xdr\_accepted\_reply, xdr\_authunix\_parms, xdr\_callhdr, xdr\_callmsg, xdr\_opaque\_auth,  
xdr\_rejected\_reply, xdr\_replymsg – XDR library routines for remote procedure calls

**DESCRIPTION**

These routines are used for describing the RPC messages in XDR language. They should normally be used by those who do not want to use the RPC package.

**Routines**

The XDR data structure is defined in the RPC/XDR Library Definitions of the *Network Programming*.

**#include <rpc/rpc.h>**

**bool\_t xdr\_accepted\_reply(xdrs, arp)**

XDR \*xdrs;

**struct accepted\_reply \*arp;**

Used for encoding RPC reply messages. It encodes the status of the RPC call in the XDR language format and in the case of success, it encodes the call results as well. This routine is useful for users who wish to generate RPC-style messages without using the RPC package. This routine returns TRUE if it succeeds, FALSE otherwise.

**bool\_t xdr\_authunix\_parms(xdrs, aup)**

XDR \*xdrs;

**struct authunix\_parms \*aup;**

Used for describing UNIX credentials. It includes machine name, user ID, group ID list, etc. This routine is useful for users who wish to generate these credentials without using the RPC authentication package. This routine returns TRUE if it succeeds, FALSE otherwise.

**void xdr\_callhdr(xdrs, chdrp)**

XDR \*xdrs;

**struct rpc\_msg \*chdrp;**

Used for describing RPC call header messages. It encodes the static part of the call message header in the XDR language format. It includes information such as transaction ID, RPC version number, program number, and version number. This routine is useful for users who wish to generate RPC-style messages without using the RPC package.

**bool\_t xdr\_callmsg(xdrs, cmsgp)**

XDR \*xdrs;

**struct rpc\_msg \*cmsgp;**

Used for describing RPC call messages. It includes all the RPC call information such as transaction ID, RPC version number, program number, version number, authentication information, etc. This routine is useful for users who wish to generate RPC-style messages without using the RPC package. This routine returns TRUE if it succeeds, FALSE otherwise.

**bool\_t xdr\_opaque\_auth(xdrs, ap)**

XDR \*xdrs;

**struct opaque\_auth \*ap;**

Used for describing RPC authentication information messages. This routine is useful for users who wish to generate RPC-style messages without using the RPC package. This routine returns TRUE if it succeeds, FALSE otherwise.

**bool\_t xdr\_rejected\_reply(xdrs, rrp)**

**XDR \*xdrs;**

**struct rejected\_reply \*rrp;**

Used for describing RPC reply messages. It encodes the rejected RPC message in the XDR language format. The message is rejected either because of version number mismatch or because of authentication errors. This routine is useful for users who wish to generate RPC-style messages without using the RPC package. This routine returns TRUE if it succeeds, FALSE otherwise.

**bool\_t xdr\_replymsg(xdrs, rmsgp)**

**XDR \*xdrs;**

**struct rpc\_msg \*rmsgp;**

Used for describing RPC reply messages. It encodes the RPC reply message in the XDR language format. This reply could be an acceptance, rejection, or NULL. This routine is useful for users who wish to generate RPC style messages without using the RPC package. This routine returns TRUE if it succeeds, FALSE otherwise.

SEE ALSO

**rpc(3N)**

**NAME**

`rtime` – get remote time

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/time.h>
#include <netinet/in.h>
```

```
int rtime(addrp, timep, timeout);
struct sockaddr_in *addrp;
struct timeval *timep;
struct timeval *timeout;
```

**DESCRIPTION**

`rtime()` consults the Internet Time Server at the address pointed to by *addrp* and returns the remote time in the `timeval` struct pointed to by *timep*. Normally, the UDP protocol is used when consulting the Time Server. The *timeout* parameter specifies how long the routine should wait before giving up when waiting for a reply. If *timeout* is specified as `NULL`, however, the routine will instead use TCP and block until a reply is received from the time server.

The routine returns 0 if it is successful. Otherwise, it returns -1 and `errno` is set to reflect the cause of the error.

**NAME**

scandir, alphasort – scan a directory

**SYNOPSIS**

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

scandir(dirname, &namelist, select, compar)
char *dirname;
struct direct **namelist;
int (*select)();
int (*compar)();

alphasort(d1, d2)
struct direct **d1, **d2;
```

**DESCRIPTION**

**scandir()** reads the directory **dirname** and builds an array of pointers to directory entries using **malloc(3V)**. The second parameter is a pointer to an array of structure pointers. The third parameter is a pointer to a routine which is called with a pointer to a directory entry and should return a non zero value if the directory entry should be included in the array. If this pointer is **NULL**, then all the directory entries will be included. The last argument is a pointer to a routine which is passed to **qsort(3)** to sort the completed array. If this pointer is **NULL**, the array is not sorted. **alphasort()** is a routine which will sort the array alphabetically.

**scandir()** returns the number of entries in the array and a pointer to the array through the parameter *namelist*.

**SEE ALSO**

**directory(3V)**, **malloc(3V)**, **qsort(3)**

**DIAGNOSTICS**

Returns **-1** if the directory cannot be opened for reading or if **malloc(3V)** cannot allocate enough memory to hold all the data structures.

## NAME

scanf, fscanf, sscanf – formatted input conversion

## SYNOPSIS

```
#include <stdio.h>

int scanf(format [ , pointer ... ])
char *format;

int fscanf(stream, format [ , pointer ... ])
FILE *stream;
char *format;

int sscanf(s, format [ , pointer ... ])
char *s, *format;
```

## SYSTEM V SYNOPSIS

The following are provided for XPG2 compatibility:

```
#define nl_scanfscanf
#define nl_fscanf      fscanf
#define nl_sscanf      sscanf
```

## DESCRIPTION

**scanf()** reads from the standard input stream **stdin**. **fscanf()** reads from the named input stream. **sscanf()** reads from the character string *s*. Each function reads characters, interprets them according to a format, and stores the results in its arguments. Each expects, as arguments, a control string *format*, described below, and a set of *pointer* arguments indicating where the converted input should be stored. The results are undefined in there are insufficient *args* for the format. If the format is exhausted while *args* remain, the excess *args* are simply ignored.

The control string usually contains conversion specifications, which are used to direct interpretation of input sequences. The control string may contain:

- White-space characters (SPACE, TAB, or NEWLINE) which, except in two cases described below, cause input to be read up to the next non-white-space character.
- An ordinary character (not '%'), which must match the next character of the input stream.
- Conversion specifications, consisting of the character '%' or the character sequence *%digit\$*, an optional assignment suppressing character '\*', an optional numerical maximum field width, an optional **l** (ell) or **h** indicating the size of the receiving variable, and a conversion code.

Conversion specifications are introduced by the character % or the character sequence *%digit\$*. A conversion specification directs the conversion of the next input field; the result is placed in the variable pointed to by the corresponding argument, unless assignment suppression was indicated by '\*'. The suppression of assignment provides a way of describing an input field which is to be skipped. An input field is defined as a string of non-space characters; it extends to the next inappropriate character or until the field width, if specified, is exhausted. For all descriptors except '[' and 'c', white space leading an input field is ignored.

The conversion character indicates the interpretation of the input field; the corresponding pointer argument must usually be of a restricted type. For a suppressed field, no pointer argument is given. The following conversion characters are legal:

- |          |  |
|----------|--|
| %        | A single % is expected in the input at this point; no assignment is done.                                  |
| <b>d</b> | A decimal integer is expected; the corresponding argument should be an integer pointer.                    |
| <b>u</b> | An unsigned decimal integer is expected; the corresponding argument should be an unsigned integer pointer. |
| <b>o</b> | An octal integer is expected; the corresponding argument should be an integer pointer.                     |
| <b>x</b> | A hexadecimal integer is expected; the corresponding argument should be an integer pointer.                |



- i** An integer is expected; the corresponding argument should be an integer pointer. It will store the value of the next input item interpreted according to C conventions: a leading “0” implies octal; a leading “0x” implies hexadecimal; otherwise, decimal.
- n** Stores in an integer argument the total number of characters (including white space) that have been scanned so far since the function call. No input is consumed.
- e,f,g** A floating point number is expected; the next field is converted accordingly and stored through the corresponding argument, which should be a pointer to a *float*. The input format for floating point numbers is as described for `string_to_decimal(3)`, with *fortran\_conventions zero*.
- s** A character string is expected; the corresponding argument should be a character pointer pointing to an array of characters large enough to accept the string and a terminating `\0`, which will be added automatically. The input field is terminated by a white space character.
- c** A character is expected; the corresponding argument should be a character pointer. The normal skip over white space is suppressed in this case; to read the next non-space character, use `%1s`. If a field width is given, the corresponding argument should refer to a character array, and the indicated number of characters is read.
- [** Indicates string data; the normal skip over leading white space is suppressed. The left bracket is followed by a set of characters, which we will call the *scanset*, and a right bracket; the input field is the maximal sequence of input characters consisting entirely of characters in the scanset. The circumflex (`^`), when it appears as the first character in the scanset, serves as a complement operator and redefines the scanset as the set of all characters *not* contained in the remainder of the scanset string. There are some conventions used in the construction of the scanset. A range of characters may be represented by the construct *first–last*, thus `[0123456789]` may be expressed `[0–9]`. Using this convention, *first* must be lexically less than or equal to *last*, or else the dash will stand for itself. The dash will also stand for itself whenever it is the first or the last character in the scanset. To include the right square bracket as an element of the scanset, it must appear as the first character (possibly preceded by a circumflex) of the scanset, and in this case it will not be syntactically interpreted as the closing bracket. The corresponding argument must point to a character array large enough to hold the data field and the terminating `\0`, which will be added automatically. At least one character must match for this conversion to be considered successful.

The conversion characters **d**, **u**, **o**, **x**, and **i** may be preceded by **l** or **h** to indicate that a pointer to **long** or to **short** rather than to **int** is in the argument list. Similarly, the conversion characters **e**, **f**, and **g** may be preceded by **l** to indicate that a pointer to **double** rather than to **float** is in the argument list. The **l** or **h** modifier is ignored for other conversion characters.

*Avoid this common error:* because `printf(3V)` does not require that the lengths of conversion descriptors and actual parameters match, coders sometimes are careless with the `scanf()` functions. But converting `%f` to `&double` or `%lf` to `&float` *does not work*; the results are quite incorrect.

`scanf()` conversion terminates at EOF, at the end of the control string, or when an input character conflicts with the control string. In the latter case, the offending character is left unread in the input stream.

`scanf()` returns the number of successfully matched and assigned input items; this number can be zero in the event of an early conflict between an input character and the control string. The constant EOF is returned upon end of input. Note: this is different from 0, which means that no conversion was done; if conversion was intended, it was frustrated by an inappropriate character in the input.

If the input ends before the first conflict or conversion, EOF is returned. If the input ends after the first conflict or conversion, the number of successfully matched items is returned.

Conversions can be applied to the  $n$ th argument in the argument list, rather than the next unused argument. In this case, the conversion character `%` (see below) is replaced by the sequence `%digit$`, where *digit* is a decimal integer  $n$  in the range [1,9], giving the position of the argument in the argument list. This feature provides for the definition of format strings that select arguments in an order appropriate to specific languages.

The format string can contain either form of a conversion specification, that is `%` or `%digit$`, although the two forms cannot be mixed within a single format string.

All forms of the `scanf()` functions allow for the detection of a language dependent radix character in the input string. The radix character is defined by the program's locale (category `LC_NUMERIC`). In the "C" locale, or in a locale where the radix character is not defined, the radix character defaults to `'.'`.

#### SYSTEM V DESCRIPTION

FORMFEED is allowed as a white space character in control strings.

XPG2 requires that `nl_scanf`, `nl_fscanf` and `nl_sscanf` be defined as `scanf`, `fscanf` and `sscanf`, respectively for backward compatibility.

#### RETURN VALUES

If any items are converted, `scanf()`, `fscanf()` and `sscanf()` return the number of items converted successfully. This number may smaller than the number of items requested. If no items are converted, these functions return 0. `scanf()`, `fscanf()` and `sscanf()` return EOF on end of input.

#### EXAMPLES

The call:

```
int i, n; float x; char name[50];
n = scanf("%d%f%s", &i, &x, name);
```

with the input line:

```
25 54.32E-1 thompson
```

will assign to  $n$  the value 3, to  $i$  the value 25, to  $x$  the value 5.432, and  $name$  will contain `thompson\0`. Or:

```
int i, j; float x; char name[50];
(void) scanf("%i%2d%f%*d %[0-9]", &j, &i, &x, name);
```

with input:

```
011 56789 0123 56a72
```

will assign 9 to  $j$ , 56 to  $i$ , 789.0 to  $x$ , skip 0123, and place the string `56\0` in  $name$ . The next call to `getchar()` (see `getc(3V)`) will return a. Or:

```
int i, j, s, e; char name[50];
(void) scanf("%i %i %n%s%n", &i, &j, &s, name, &e);
```

with input:

```
0x11 0xy johnson
```

will assign 17 to  $i$ , 0 to  $j$ , 6 to  $s$ , will place the string `xy\0` in  $name$ , and will assign 8 to  $e$ . Thus, the length of  $name$  is  $e - s = 2$ . The next call to `getchar()` (see `getc(3V)`) will return a SPACE.

#### SEE ALSO

`getc(3V)`, `printf(3V)`, `setlocale(3V)`, `stdio(3V)`, `string_to_decimal(3)`, `strtol(3)`

**WARNINGS**

Trailing white space (including a NEWLINE) is left unread unless matched in the control string.

**BUGS**

The success of literal matches and suppressed assignments is not directly determinable.

## NAME

authdes\_create, authdes\_getucred, get\_myaddress, getnetname, host2netname, key\_decryptsession, key\_encryptsession, key\_gendes, key\_setsecret, netname2host, netname2user, user2netname – library routines for secure remote procedure calls

## DESCRIPTION

RPC routines allow C programs to make procedure calls on other machines across the network. First, the client calls a procedure to send a request to the server. Upon receipt of the request, the server calls a dispatch routine to perform the requested service, and then sends back a reply. Finally, the procedure call returns to the client.

RPC allows various authentication flavors. The `authdes_getucred()` and `authdes_create()` routines implement the DES authentication flavor. See `rpc_clnt_auth(3N)` for routines relating to the `AUTH_NONE` and `AUTH_UNIX` authentication types.

Note: Both the client and server should have their keys in the `publickey(5)` database. Also, the keyserver daemon `keyerv(8C)` must be running on both the client and server hosts for the DES authentication system to work.

## Routines

```
#include <rpc/rpc.h>
```

```
AUTH * authdes_create(netname, window, syncaddr, deskeyp)
```

```
char *netname;
```

```
unsigned window;
```

```
struct sockaddr_in *syncaddr;
```

```
des_block *deskeyp;
```

`authdes_create()` is an interface to the RPC secure authentication system, known as DES authentication.

Used on the client side, `authdes_create()` returns an authentication handle that enables the use of the secure authentication system. The first parameter *netname* is the network name of the owner of the server process. This field usually represents a *host* derived from the utility routine `host2netname()`, but could also represent a user name using `user2netname()`. The second field is *window* on the validity of the client credential, given in seconds. A small window is more secure than a large one, but choosing too small of a window will increase the frequency of resynchronizations because of clock drift. The third parameter *syncaddr* is optional. If it is `NULL`, then the authentication system will assume that the local clock is always in sync with the server's clock, and will not attempt to synchronize with the server. If an address is supplied then the system will use it for consulting the remote time service whenever resynchronization is required. This parameter is usually the address of the RPC server itself. The final parameter *deskeyp* is also optional. If it is `NULL`, then the authentication system will generate a random DES key to be used for the encryption of credentials. If *deskeyp* is supplied then it is used instead.

```
int authdes_getucred(adc, uidp, gidp, gidlenp, gidlistp)
```

```
struct authdes_cred *adc;
```

```
short *uidp;
```

```
short *gidp;
```

```
short *gidlenp;
```

```
int *gidlistp;
```

`authdes_getucred()`, is a DES authentication routine used by the server for converting a DES credential, which is operating system independent, into a UNIX credential. *uidp* points to the user ID of the user associated with *adc*; *gidp* refers to the user's current group ID; *gidlistp* refers to an array of groups to which the user belongs and *gidlenp* has the count of the entries in this array.

This routine differs from the utility routine `netname2user()` in that `authdes_getucred()` pulls its information from a cache, and does not have to do a NIS name service lookup every time it is called to get its information. Returns 1 if it succeeds and 0 if it fails.

```
void get_myaddress(addr)
struct sockaddr_in *addr;
```

Return the machine's IP address in `addr`. The port number is always set to `htons(PMAPPORT)`.

```
int getnetname(netname)
char netname[MAXNETNAMELEN];
```

Return the unique, operating-system independent netname of the caller in the fixed-length array `netname`. Returns 1 if it succeeds and 0 if it fails.

```
int host2netname(netname, host, domain)
char netname[MAXNETNAMELEN];
char *host;
char *domain;
```

Convert from a domain-specific hostname to an operating-system independent netname. This routine is normally used to get the netname of the server, which is then used to get an authentication handle by calling `authdes_create()`. This routine should be used if the owner of the server process is the machine that is, the user with effective user ID zero. Returns 1 if it succeeds and 0 if it fails. This routine is the inverse of `netname2host()`.

```
int key_decryptsession(netname, deskeyp)
char *netname;
des_block *deskeyp;
```

An interface routine to the keyserver daemon, which is associated with RPC's secure authentication system (DES authentication). User programs rarely need to call it, or its associated routines `key_encryptsession()`, `key_gendes()` and `key_setsecret()`. System commands such as `login` and the RPC library are the main clients of these four routines.

`key_decryptsession()` takes the netname of a server and a DES key, and decrypts the key by using the public key of the server and the secret key associated with the effective user ID of the calling process. Returns 0 if it succeeds and -1 if it fails. This routine is the inverse of `key_encryptsession()`.

```
int key_encryptsession(netname, deskeyp)
char *netname;
des_block *deskeyp;
```

A keyserver interface routine. It takes the netname of the server and a des key, and encrypts it using the public key of the server and the secret key associated with the effective user ID of the calling process. Returns 0 if it succeeds and -1 if it fails. This routine is the inverse of `key_decryptsession()`.

```
int key_gendes(deskeyp)
des_block *deskeyp;
```

A keyserver interface routine. It is used to ask the keyserver for a secure conversation key. Choosing one at "random" is usually not good enough, because the common ways of choosing random numbers, such as using the current time, are very easy to guess. Returns 0 if it succeeds and -1 if it fails.

```
int key_setsecret(keyp)
char *keyp;
```

A keyserver interface routine. It is used to set the secret key for the effective user ID of the calling process. Returns 0 if it succeeds and -1 if it fails.

```
int netname2host(netname, host, hostlen)
char *netname;
char *host;
int hostlen;
```

Convert an operating-system independent netname to a domain-specific hostname. *hostlen* specifies the size of the array pointed to by *host*. It returns 1 if it succeeds and 0 if it fails. This routine is the inverse of `host2netname()`.

```
int netname2user(netname, uidp, gidp, gidlenp, gidlistp)
char *name;
int *uidp;
int *gidp;
int *gidlenp;
int *gidlistp;
```

Convert an operating-system independent netname to a domain-specific user ID. *uidp* points to the user ID of the user; *gidp* refers to the user's current group ID; *gidlistp* refers to an array of groups to which the user belongs and *gidlenp* has the count of the entries in this array. It returns 1 if it succeeds and 0 if it fails. This routine is the inverse of `user2netname()`.

```
int user2netname(netname, uid, domain)
char name[MAXNETNAMELEN];
int uid;
char *domain;
```

Convert a domain-specific username to an operating-system independent netname. *uid* is the user ID of the owner of the server process. This routine is normally used to get the netname of the server, which is then used to get an authentication handle by calling `authdes_create()`. Returns 1 if it succeeds and 0 if it fails. This routine is the inverse of `netname2user()`.

#### SEE ALSO

`login(1)`, `chkey(1)`, `rpc(3N)`, `rpc_clnt_auth(3N)`, `publickey(5)`, `keyserv(8C)`, `newkey(8)`

## NAME

setbuf, setbuffer, setlinebuf, setvbuf – assign buffering to a stream

## SYNOPSIS

```
#include <stdio.h>

void setbuf(stream, buf)
FILE *stream;
char *buf;

void setbuffer(stream, buf, size)
FILE *stream;
char *buf;
int size;

int setlinebuf(stream) FILE *stream;

int setvbuf(stream, buf, type, size)
FILE *stream;
char *buf;
int type, size;
```

## DESCRIPTION

The three types of buffering available are unbuffered, block buffered, and line buffered. When an output stream is unbuffered, information appears on the destination file or terminal as soon as written; when it is block buffered many characters are saved up and written as a block; when it is line buffered characters are saved up until a NEWLINE is encountered or input is read from `stdin`. `fflush()` (see `fclose(3V)`) may be used to force the block out early. A buffer is obtained from `malloc(3V)` upon the first `getc(3V)` or `putc(3S)` on the file. By default, output to a terminal is line buffered, except for output to the standard stream `stderr` which is unbuffered. All other input/output is fully buffered.

`setbuf()` can be used after a stream has been opened but before it is read or written. It causes the array pointed to by `buf` to be used instead of an automatically allocated buffer. If `buf` is the NULL pointer, input/output will be completely unbuffered. A manifest constant `BUFSIZ`, defined in the `<stdio.h>` header file, tells how big an array is needed:

```
char buf[BUFSIZ];
```

`setbuffer()`, an alternate form of `setbuf()`, can be used after a stream has been opened but before it is read or written. It uses the character array `buf` whose size is determined by the `size` argument instead of an automatically allocated buffer. If `buf` is the NULL pointer, input/output will be completely unbuffered.

`setvbuf()` can be used after a stream has been opened but before it is read or written. `type` determines how stream will be buffered. Legal values for `type` (defined in `<stdio.h>`) are:

```
_IOFBF    fully buffers the input/output.
_IOLBF    line buffers the output; the buffer will be flushed when a NEWLINE is written, the
           buffer is full, or input is requested.
_IONBF    completely unbuffers the input/output.
```

If `buf` is not the NULL pointer, the array it points to will be used for buffering, instead of an automatically allocated buffer. `size` specifies the size of the buffer to be used.

`setlinebuf()` is used to change the buffering on a stream from block buffered or unbuffered to line buffered. Unlike `setbuf()`, `setbuffer()`, and `setvbuf()`, it can be used at any time that the file descriptor is active.

A file can be changed from unbuffered or line buffered to block buffered by using `freopen()` (see `fopen(3V)`). A file can be changed from block buffered or line buffered to unbuffered by using `freopen()` followed by `setbuf()` with a buffer argument of `NULL`.

**SYSTEM V DESCRIPTION**

If *buf* is not `NULL` and *stream* refers to a terminal device, `setbuf()` sets *stream* for line buffered input/output.

**RETURN VALUES**

`setlinebuf()` returns no useful value.

`setvbuf()` returns 0 on success. If an illegal value for *type* or *size* is provided, `setvbuf()` returns a non-zero value. `setvbuf()`

**SEE ALSO**

`fclose(3V)`, `fopen(3V)`, `fread(3S)`, `getc(3V)`, `malloc(3V)`, `printf(3V)`, `putc(3S)`, `puts(3S)`

**NOTES**

A common source of error is allocating buffer space as an "automatic" variable in a code block, and then failing to close the stream in the same block.



## NAME

setjmp, longjmp, sigsetjmp, siglongjmp – non-local goto

## SYNOPSIS

```
#include <setjmp.h>

int setjmp(env)
jmp_buf env;

void longjmp(env, val)
jmp_buf env;
int val;

int _setjmp(env)
jmp_buf env;

void _longjmp(env, val)
jmp_buf env;
int val;

int sigsetjmp(env, savemask)
sigjmp_buf env;
int savemask;

void siglongjmp(env, val)
sigjmp_buf env;
int val;
```

## DESCRIPTION

**setjmp()** and **longjmp()** are useful for dealing with errors and interrupts encountered in a low-level subroutine of a program.

The macro **setjmp()** saves its stack environment in *env* for later use by **longjmp()**. A normal call to **setjmp()** returns zero. **setjmp()** also saves the register environment. If a **longjmp()** call will be made, the routine which called **setjmp()** should not return until after the **longjmp()** has returned control (see below).

**longjmp()** restores the environment saved by the last call of **setjmp**, and then returns in such a way that execution continues as if the call of **setjmp()** had just returned the value *val* to the function that invoked **setjmp()**; however, if *val* were zero, execution would continue as if the call of **setjmp()** had returned one. This ensures that a “return” from **setjmp()** caused by a call to **longjmp()** can be distinguished from a regular return from **setjmp()**. The calling function must not itself have returned in the interim, otherwise **longjmp()** will be returning control to a possibly non-existent environment. All memory-bound data have values as of the time **longjmp()** was called. The CPU and floating-point data registers are restored to the values they had at the time that **setjmp()** was called. But, because the **register** storage class is only a hint to the C compiler, variables declared as **register** variables may not necessarily be assigned to machine registers, so their values are unpredictable after a **longjmp()**. This is especially a problem for programmers trying to write machine-independent C routines.

**setjmp()** and **longjmp()** save and restore the signal mask (see **sigsetmask(2)**), while **\_setjmp()** and **\_longjmp()** manipulate only the C stack and registers. If the *savemask* flag to **sigsetjmp()** is non-zero, the signal mask is saved, and a subsequent **siglongjmp()** using the same *env* will restore the signal mask. If the *savemask* flag is zero, the signal mask is not saved, and a subsequent **siglongjmp()** using the same *env* will not restore the signal mask. In all other ways, **\_setjmp()** and **sigsetjmp()** function in the same way that **setjmp()** does, and **\_longjmp()** and **siglongjmp()** function in the same way that **longjmp()** does.

None of these functions save or restore any floating-point status or control registers, in particular the MC68881 **fpsr**, **fpcr**, or **fpiar**, the Sun-3 FPA **fpamode** or **fpastatus**, and the Sun-4 **%fsr**. See **ieee\_flags(3M)** to save and restore floating-point status or control information.

**SYSTEM V DESCRIPTION**

**setjmp()** and **longjmp()** manipulate only the C stack and registers; they do not save or restore the signal mask. **\_setjmp()** behaves identically to **setjmp()**, and **\_longjmp()** behaves identically to **longjmp()**.

**EXAMPLE**

The following code fragment indicates the flow of control of the **setjmp()** and **longjmp()** combination:

```

function declaration
...
    jmp_buf    my_environment;
    ...
    if (setjmp(my_environment)) {
        /* register variables have unpredictable values */
        /* code after the return from longjmp */
        ...
    } else {
        /* do not modify register vars in this leg of code */
        /* this is the return from setjmp */
        ...
    }

```

**SEE ALSO**

**cc(1V)**, **sigsetmask(2)**, **sigvec(2)**, **ieee\_flags(3M)**, **signal(3V)**, **setjmp(3V)**

**BUGS**

**setjmp()** does not save the current notion of whether the process is executing on the signal stack. The result is that a **longjmp()** to some place on the signal stack leaves the signal stack state incorrect.

On Sun-2 and Sun-3 systems **setjmp()** also saves the register environment. Therefore, all data that are bound to registers are restored to the values they had at the time that **setjmp()** was called. All memory-bound data have values as of the time **longjmp()** was called. However, because the **register** storage class is only a hint to the C compiler, variables declared as **register** variables may not necessarily be assigned to machine registers, so their values are unpredictable after a **longjmp()**. When using compiler options that specify automatic register allocation (see **cc(1V)**), the compiler will not attempt to assign variables to registers in routines that call **setjmp()**.

**NAME**

setlocale, nl\_init – set international environment

**SYNOPSIS**

```
#include <locale.h>
```

```
char *setlocale(category, locale)
```

```
int category;
```

```
char *locale;
```

```
int nl_init(lang)
```

```
char *lang;
```

**DESCRIPTION**

**setlocale()** selects the appropriate piece of the program's locale as specified by *category*, and may be used to change or query the program's international environment. The entire locale may be changed by calling **setlocale()** with *category* set to **LC\_ALL**. The other possible values for *category* query or change only a part of the program's complete international locale:

**LC\_CTYPE**

Affects the behavior of the character classification and conversion functions. See **ctype(3V)**, and **mblen(3)**.

**LC\_COLLATE**

Affects the behavior of the string collation functions **strcoll(3)** and **strxfrm(3V)**.

**LC\_TIME**

Affects the behavior of the time conversion functions. See **printf(3V)**, **scanf(3V)**, **strtod(3)**, and **ctime(3V)** for **strftime()**, **strptime()**, and **ctime()**.

**LC\_NUMERIC**

Affects the radix character for the formatted input/output functions and the string conversion functions, **gcvt(3V)**, **printf(3V)**, **strtod(3)**, **gconvert()**, **sgconvert()** (see **econvert(3)**), **file\_to\_decimal()**, and **func\_to\_decimal()** (see **string\_to\_decimal(3)**). Also affects the non-monetary formatting information returned by the **localeconv()** function.

**LC\_MONETARY**

Affects the monetary formatting information returned by the **localeconv()** function.

**LC\_MESSAGES**

Affects the behavior of functions that present messages, namely **gettext()**, and **textdomain()**.

The *locale* argument is a pointer to a character string containing the required setting of *category*. The following preset values of *locale* are defined for all settings of *category*:

"C" Specifies the minimal environment for C translation. If **setlocale()** is not invoked, the "C" locale is the default. Operational behavior within the "C" locale is defined separately for each interface function.

At program startup, the equivalent of:

```
""
```

In this case, **setlocale()** will first check the value of the corresponding environment variable (for example, **LC\_CTYPE** for the **LC\_CTYPE** category) and if valid (that is, points to the name of a valid locale), **setlocale()** sets the specified category of the international environment to that value and returns the string corresponding to the locale set (that is, the value of the environment variable, not ""). If the value is invalid, **setlocale()** returns a NULL pointer and the international environment is not changed by this call.

If the environment variable corresponding to the specified category is not set or is set to the empty string, **setlocale()** will examine the **LANG** environment variable. If both the **LANG** environment variable, and the environment variable corresponding to the specified category are not set or are set to the empty string, then the **LC\_default** environment variable is examined. If this contains a valid setting, then the category is set to the value of **LC\_default**. If

the LANG environment variable is set and valid this will set the category to the corresponding value of LANG. If LC\_default is not set, then setlocale() returns that category to the default "C" locale.

To set all categories in the international environment, setlocale() is invoked in the following manner:

```
setlocale(LC_ALL, "");
```

To satisfy this request, setlocale() first checks all the relevant environment variables LC\_CTYPE, LC\_COLLATE, LC\_TIME, LC\_NUMERIC, LC\_MONETARY, LC\_MESSAGES. If any one of these relevant environment variables is invalid, this call to setlocale() will return a NULL pointer, and the international environment will not be changed. If all the relevant environment variables are valid, setlocale() sets the international environment to reflect the values of the environment variables. The categories are set in the following order:

```
LC_CTYPE
LC_COLLATE
LC_TIME
LC_NUMERIC
LC_MONETARY
LC_MESSAGES
```

Using this scheme, the categories corresponding to the environment variables will override the value of the LANG and LC\_default environment variables for a particular category.

nl\_init() is equivalent to

```
setlocale(LC_ALL, "");
```

and is supplied for compatibility with X/Open XPG2.

#### RETURN VALUES

If a valid string is given for the *locale* parameter, and the selection can be honored, setlocale() returns the string associated with the specified *category* for the new locale. If the selection cannot be honored, setlocale() returns a null pointer and the program's locale is not changed.

A NULL pointer for *locale* causes setlocale() to return the string associated with the *category* for the program's current locale; the program's locale is not changed. The string contains information relating to each piece part of the whole international environment. This inquiry can fail by returning a null pointer if any *category* is invalid.

The string returned by such a setlocale() call is such that a subsequent call with the string and its associated category will restore that part of the program's locale. The string returned by:

```
ptr = setlocale(LC_ALL, (char *) 0);
```

is such that in a subsequent call:

```
setlocale(LC_ALL, ptr);
```

will reset each and every category to the state when the string was first returned. The string returned must not be modified by the program, but will be overwritten by a subsequent call to setlocale().

#### FILES

*/etc/locale/locale/category*

*locale* is the directory that contains numerous files (*categories*), each relating to a single category of a valid *locale* as selected by category argument to setlocale(). Generally this is classed as a private directory. This directory is searched by setlocale(), prior to searching:

*/usr/share/lib/locale/locale/category*

*locale* is the directory that contains numerous files (*categories*), each relating to a single category of a valid *locale* as selected by category argument to setlocale(). Generally this data is classed as global and sharable.

**DIAGNOSTICS**

**setlocale()** returns a null pointer if a relevant environment variable has an invalid setting. **setlocale()** also returns a null pointer if *category* is invalid.

**NAME**

setuid, seteuid, setruid, setgid, setegid, setrgid – set user and group ID

**SYNOPSIS**

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

```
int setuid(uid)
```

```
uid_t uid;
```

```
int seteuid(euid)
```

```
uid_t euid;
```

```
int setruid(ruid)
```

```
uid_t ruid;
```

```
int setgid(gid)
```

```
gid_t gid;
```

```
int setegid(egid)
```

```
gid_t egid;
```

```
int setrgid(rgid)
```

```
gid_t rgid;
```

**DESCRIPTION**

**setuid()** (**setgid()**) sets both the real and effective user ID (group ID) of the current process as specified by *uid* (*gid*) (see NOTES).

**seteuid()** (**setegid()**) sets the effective user ID (group ID) of the current process.

**setruid()** (**setrgid()**) sets the real user ID (group ID) of the current process.

These calls are only permitted to the super-user or if the argument is the real or effective user (group) ID of the calling process.

**SYSTEM V DESCRIPTION**

If the effective user ID of the calling process is not super-user, but if its real user (group) ID is equal to *uid* (*gid*), or if the saved set-user (group) ID from **execve(2V)** is equal to *uid* (*gid*), then the effective user (group) ID is set to *uid* (*gid*).

**RETURN VALUES**

These functions return:

0        on success.

-1        on failure and set **errno** to indicate the error as for **setreuid(2)** (**setregid(2)**).

**ERRORS**

**EINVAL**        The value of *uid* (*gid*) is invalid (less than 0 or greater than 65535).

**EPERM**        The process does not have super-user privileges and *uid* (*gid*) does not match either the real user (group) ID of the process nor the saved set-user-ID (set-group-ID) of the process.

**SEE ALSO**

**execve(2V)**, **getgid(2V)**, **getuid(2V)**, **setregid(2)**, **setreuid(2)**

**NOTES**

For **setuid()** to behave as described above, **{\_POSIX\_SAVED\_IDS}** must be in effect (see **sysconf(2V)**). **{\_POSIX\_SAVED\_IDS}** is always in effect on SunOS systems, but for portability, applications should call **sysconf()** to determine whether **{\_POSIX\_SAVED\_IDS}** is in effect for the current system.

## NAME

sigaction – examine and change signal action

## SYNOPSIS

```
#include <signal.h>

int sigaction(sig, act, oact)
int sig;
struct sigaction *act, *oact;
```

## DESCRIPTION

**sigaction()** allows the calling process to examine and specify (or both) the action to be associated with a specific signal. *sig* specifies the signal. Acceptable values are defined in **<signal.h>**.

The structure **sigaction()**, used to describe an action to be taken, is defined in the header **<signal.h>** as follows:

```
struct sigaction {
    void (*sa_handler)(); /* SIG_DFL, SIG_IGN, or pointer to a function */
    sigset_t sa_mask; /* Additional signals to be blocked during
                       execution of signal-catching function */
    int sa_flags; /* Special flags to affect behavior of signal */
};
```

If *act* is not NULL, it points to a structure specifying the action to be associated with the specified signal. If *oact* is not NULL, the action previously associated with the signal is stored in the location pointed to by the *oact*. If *act* is NULL, signal handling is unchanged by this function. Thus, the call can be used to enquire about the current handling of a given signal. The **sa\_handler** field of the **sigaction** structure identifies the action to be associated with the specified signal. If the **sa\_handler** field specifies a signal-catching function, the **sa\_mask** field identifies a set of signals that shall be added to the process's signal mask before the signal-catching function mask is invoked. The SIGKILL and SIGSTOP signals shall not be added to the signal mask using this mechanism; this restriction shall be enforced by the system without causing an error to be indicated.

The **sa\_flags** field can be used to modify the behavior of the specified signal. The following flag bit, defined in the header **<signal.h>**, can be set in **sa\_flags**:

```
#define SA_ONSTACK      0x0001 /* take signal on signal stack */
#define SA_INTERRUPT    0x0002 /* do not restart system on signal return */
#define SA_RESETHAND    0x0004 /* reset handler to SIG_DFL when signal taken */
#define SA_NOCLDSTOP    0x0008 /* don't send a SIGCHLD on child stop */
```

If *sig* is SIGCHLD and the SA\_NOCLDSTOP flag is not set in **sa\_flags**, and the implementation supports the SIGCHLD signal, a SIGCHLD signal shall be generated for the calling process whenever any of its child processes stop. If *sig* is SIGCHLD and the SA\_NOCLDSTOP flag is set in **sa\_flags**, the implementation shall not generate a SIGCHLD signal in this way.

If the SA\_ONSTACK bit is set in the flags for that signal, the system will deliver the signal to the process on the signal stack specified with **sigstack(2)**, rather than delivering the signal on the current stack.

If a caught signal occurs during certain system calls, the call is restarted by default. The call can be forced to terminate prematurely with an EINTR error return by setting the SA\_INTERRUPT bit in the flags for that signal. SA\_INTERRUPT is not available in 4.2BSD, hence it should not be used if backward compatibility is needed. The affected system calls are **read(2V)** or **write(2V)** on a slow device (such as a terminal or pipe or other socket, but not a file) and during a **wait(2V)**.

Once a signal handler is installed, it remains installed until another **sigvec()** call is made, or an **execve(2V)** is performed, unless the SA\_RESETHAND bit is set in the flags for that signal. In that case, the value of the handler for the caught signal is set to SIG\_DFL before entering the signal-catching function, unless the signal is SIGILL or SIGTRAP. Also, if this bit is set, the bit for that

signal in the signal mask will not be set; unless the signal mask associated with that signal blocks that signal, further occurrences of that signal will not be blocked. The SA\_RESETHAND flag is not available in 4.2BSD, hence it should not be used if backward compatibility is needed.

When a signal is caught by a signal-catching function installed by `sigaction()` a new signal mask is calculated and installed for the duration of the signal-catching function (or until a call to either `sigprocmask()` or `sigsuspend()`). This mask is formed by taking the union of the current signal mask and the value of the `sa_mask` for the signal being delivered, and then including the signal being delivered. If and when the user's signal handler returns normally, the original signal mask is restored.

Once an action is installed for a specific signal, it remains installed until another action is explicitly requested (by another call to `sigaction()`), or until one of the `exec` functions is called.

If the previous action for `sig` had been established by `signal()` defined in the C standard, the values of the fields returned in the structure pointed to by the `oact` are unspecified, and in particular `oact->sv_handler` is not necessarily the same value passed to `signal()`. However, if a pointer to the same structure or a copy thereof is passed to a subsequent call to `sigaction()` using `act`, handling of the signal shall be as if the original call to `signal()` were repeated.

If `sigaction()` fails, no new signal handler is installed.

#### RETURN VALUES

`sigaction()` returns:

- 0        on success.
- 1       on failure and sets `errno` to indicate the error.

#### ERRORS

- EINVAL        `sig` is an invalid or unsupported signal number.  
An attempt was made to catch a signal that cannot be ignored. See `<signal.h>`.

#### SEE ALSO

`kill(2V)`, `sigpause(2V)`, `sigprocmask(2V)`, `signal(3V)`, `sigsetops(3V)`



**NAME**

`sigfpe` – signal handling for specific SIGFPE codes

**SYNOPSIS**

```
#include <signal.h>
#include <floatingpoint.h>

sigfpe_handler_type sigfpe(code, hdl)
sigfpe_code_type code;
sigfpe_handler_type hdl;
```

**DESCRIPTION**

This function allows signal handling to be specified for particular SIGFPE codes. A call to `sigfpe()` defines a new handler *hdl* for a particular SIGFPE *code* and returns the old handler as the value of the function `sigfpe()`. Normally handlers are specified as pointers to functions; the special cases SIGFPE\_IGNORE, SIGFPE\_ABORT, and SIGFPE\_DEFAULT allow ignoring, specifying core dump using `abort(3)`, or default handling respectively.

For these IEEE-related codes:

FPE_FLTINEX_TRAP	fp_inexact - floating inexact result
FPE_FLTDIV_TRAP	fp_division - floating division by zero
FPE_FLTUND_TRAP	fp_underflow - floating underflow
FPE_FLTOVF_TRAP	fp_overflow - floating overflow
FPE_FLTBSUN_TRAP	fp_invalid - branch or set on unordered
FPE_FLTOPERR_TRAP	fp_invalid - floating operand error
FPE_FLTNAN_TRAP	fp_invalid - floating Not-A-Number

default handling is defined to be to call the handler specified to `ieee_handler(3M)`.

For all other SIGFPE codes, default handling is to core dump using `abort(3)`.

The compilation option `-ffpa` causes fpa recomputation to replace the default abort action for code FPE\_FPA\_ERROR. Note: SIGFPE\_DEFAULT will restore abort rather than FPA recomputation for this code.

Three steps are required to intercept an IEEE-related SIGFPE code with `sigfpe()`:

- 1) Set up a handler with `sigfpe()`.
- 2) Enable the relevant IEEE trapping capability in the hardware, perhaps by using assembly-language instructions.
- 3) Perform a floating-point operation that generates the intended IEEE exception.

Unlike `ieee_handler(3M)`, `sigfpe()` never changes floating-point hardware mode bits affecting IEEE trapping. No IEEE-related SIGFPE signals will be generated unless those hardware mode bits are enabled.

SIGFPE signals can be handled using `sigvec(2)`, `signal(3V)`, `sigfpe(3)`, or `ieee_handler(3M)`. In a particular program, to avoid confusion, use only one of these interfaces to handle SIGFPE signals.

## EXAMPLE

A user-specified signal handler might look like this:

```
void sample_handler( sig, code, scp, addr )
    int sig ;          /* sig == SIGFPE always */
    int code ;
    struct sigcontext *scp ;
    char *addr ;
    {
        /*
         * Sample user-written sigfpe code handler.
         * Prints a message and continues.
         * struct sigcontext is defined in <signal.h>.
         */
        printf(" ieee exception code %x occurred at pc %X \n",code,scp->sc_pc);
    }
```

and it might be set up like this:

```
extern void sample_handler();
main()
{
    sigfpe_handler_type hdl, old_handler1, old_handler2;
    /*
     * save current overflow and invalid handlers; set the new
     * overflow handler to sample_handler() and set the new
     * invalid handler to SIGFPE_ABORT (abort on invalid)
     */
    hdl = (sigfpe_handler_type) sample_handler;
    old_handler1 = sigfpe(FPE_FLTOVF_TRAP, hdl);
    old_handler2 = sigfpe(FPE_FLTOPERR_TRAP, SIGFPE_ABORT);
    ...
    /*
     * restore old overflow and invalid handlers
     */
    sigfpe(FPE_FLTOVF_TRAP, old_handler1);
    sigfpe(FPE_FLTOPERR_TRAP, old_handler2);
}
```

## SEE ALSO

sigvec(2), abort(3), floatingpoint(3), ieee\_handler(3M), signal(3V)

## DIAGNOSTICS

sigfpe() returns BADSIG if *code* is not zero or a defined SIGFPE code.

**NAME**

`siginterrupt` – allow signals to interrupt system calls

**SYNOPSIS**

```
int siginterrupt(sig, flag)  
int sig, flag;
```

**DESCRIPTION**

`siginterrupt()` is used to change the system call restart behavior when a system call is interrupted by the specified signal. If the flag is false (0), then system calls will be restarted if they are interrupted by the specified signal and no data has been transferred yet. System call restart is the default behavior on 4.2BSD, and on SunOS in the 4.2 environment, when the `signal(3V)` routine is used.

If the flag is true (1), then restarting of system calls is disabled. If a system call is interrupted by the specified signal and no data has been transferred, the system call will return -1 with `errno` set to `EINTR`. Interrupted system calls that have started transferring data will return the amount of data actually transferred. System call interrupt is the signal behavior found on older version of the UNIX operating systems, such as 4.1BSD and System V UNIX. It is the default behavior on SunOS in the System V environment when the `signal()` routine is used; therefore, this routine is useful in that environment only if a signal that a `sigvec(2)` specified should restart system calls is to be changed not to restart them.

Note: the new 4.2BSD signal handling semantics are not altered in any other way. Most notably, signal handlers always remain installed until explicitly changed by a subsequent `sigvec()` call, and the signal mask operates as documented in `sigvec()`, unless the `SV_RESETHAND` bit has been used to specify that the pre-4.2BSD signal behavior is to be used. Programs may switch between restartable and interruptible system call operation as often as desired in the execution of a program.

Issuing a `siginterrupt()` call during the execution of a signal handler will cause the new action to take place on the next signal to be caught.

**NOTES**

This library routine uses an extension of the `sigvec(2)` system call that is not available in 4.2BSD, hence it should not be used if backward compatibility is needed.

**RETURN VALUES**

`siginterrupt()` returns:

- 0       on success.
- 1       if an invalid signal number was supplied.

**SEE ALSO**

`sigblock(2)`, `sigpause(2V)`, `sigsetmask(2)`, `sigvec(2)`, `signal(3V)`

## NAME

signal – simplified software signal facilities

## SYNOPSIS

```
#include <signal.h>

void (*signal(sig, func))()
void (*func)();
```

## DESCRIPTION

**signal()** is a simplified interface to the more general **sigvec(2)** facility. Programs that use **signal()** in preference to **sigvec()** are more likely to be portable to all systems.

A signal is generated by some abnormal event, initiated by a user at a terminal (quit, interrupt, stop), by a program error (bus error, etc.), by request of another program (kill), or when a process is stopped because it wishes to access its control terminal while in the background (see **termio(4)**). Signals are optionally generated when a process resumes after being stopped, when the status of child processes changes, or when input is ready at the control terminal. Most signals cause termination of the receiving process if no action is taken; some signals instead cause the process receiving them to be stopped, or are simply discarded if the process has not requested otherwise. Except for the **SIGKILL** and **SIGSTOP** signals, the **signal()** call allows signals either to be ignored or to interrupt to a specified location. The following is a list of all signals with names as in the include file **<signal.h>**:

<b>SIGHUP</b>	1	hangup
<b>SIGINT</b>	2	interrupt
<b>SIGQUIT</b>	3*	quit
<b>SIGILL</b>	4*	illegal instruction
<b>SIGTRAP</b>	5*	trace trap
<b>SIGABRT</b>	6*	abort (generated by <b>abort(3)</b> routine)
<b>SIGEMT</b>	7*	emulator trap
<b>SIGFPE</b>	8*	arithmetic exception
<b>SIGKILL</b>	9	kill (cannot be caught, blocked, or ignored)
<b>SIGBUS</b>	10*	bus error
<b>SIGSEGV</b>	11*	segmentation violation
<b>SIGSYS</b>	12*	bad argument to system call
<b>SIGPIPE</b>	13	write on a pipe or other socket with no one to read it
<b>SIGALRM</b>	14	alarm clock
<b>SIGTERM</b>	15	software termination signal
<b>SIGURG</b>	16●	urgent condition present on socket
<b>SIGSTOP</b>	17†	stop (cannot be caught, blocked, or ignored)
<b>SIGTSTP</b>	18†	stop signal generated from keyboard
<b>SIGCONT</b>	19●	continue after stop
<b>SIGCHLD</b>	20●	child status has changed
<b>SIGTTIN</b>	21†	background read attempted from control terminal
<b>SIGTTOU</b>	22†	background write attempted to control terminal
<b>SIGIO</b>	23●	I/O is possible on a descriptor (see <b>fcntl(2V)</b> )
<b>SIGXCPU</b>	24	cpu time limit exceeded (see <b>getrlimit(2)</b> )
<b>SIGXFSZ</b>	25	file size limit exceeded (see <b>getrlimit(2)</b> )
<b>SIGVTALRM</b>	26	virtual time alarm (see <b>getitimer(2)</b> )
<b>SIGPROF</b>	27	profiling timer alarm (see <b>getitimer(2)</b> )
<b>SIGWINCH</b>	28●	window changed (see <b>termio(4)</b> and <b>win(4S)</b> )
<b>SIGLOST</b>	29*	resource lost (see <b>lockd(8C)</b> )
<b>SIGUSR1</b>	30	user-defined signal 1
<b>SIGUSR2</b>	31	user-defined signal 2

The starred signals in the list above cause a core image if not caught or ignored.

If *func* is SIG\_DFL, the default action for signal *sig* is reinstated; this default is termination (with a core image for starred signals) except for signals marked with ● or †. Signals marked with ● are discarded if the action is SIG\_DFL; signals marked with † cause the process to stop. If *func* is SIG\_IGN the signal is subsequently ignored and pending instances of the signal are discarded. Otherwise, when the signal occurs further occurrences of the signal are automatically blocked and *func* is called.

A return from the function unblocks the handled signal and continues the process at the point it was interrupted. Unlike previous signal facilities, the handler *func* remains installed after a signal has been delivered.

If a caught signal occurs during certain system calls, terminating the call prematurely, the call is automatically restarted. In particular this can occur during a read(2V) or write(2V) on a slow device (such as a terminal; but not a file) and during a wait(2V).

The value of signal() is the previous (or initial) value of *func* for the particular signal.

After a fork(2V) or vfork(2) the child inherits all signals. An execve(2V) resets all caught signals to the default action; ignored signals remain ignored.

#### SYSTEM V DESCRIPTION

If *func* is SIG\_IGN the signal is subsequently ignored and pending instances of the signal are discarded. Otherwise, when the signal occurs, *func* is called. Further occurrences of the signal are not automatically blocked. The value of *func* for the caught signal is reset to SIG\_DFL before *func* is called, unless the signal is SIGILL or SIGTRAP.

A return from the function continues the process at the point at which it was interrupted. The handler *func* does not remain installed after a signal has been delivered.

If a caught signal occurs during certain system calls, causing the call to terminate prematurely, the call is interrupted. In particular this can occur during a read(2V) or write(2V) on a slow device (such as a terminal; but not a file) and during a wait(2V). After the signal catching function returns, the interrupted system call may return a -1 to the calling process with errno set to EINTR.

#### RETURN VALUES

signal() returns the previous action on success. On failure, it returns -1 and sets errno to indicate the error.

#### ERRORS

signal() will fail and no action will take place if one of the following occurs:

EINVAL            *sig* was not a valid signal number.

An attempt was made to ignore or supply a handler for SIGKILL or SIGSTOP.

#### SEE ALSO

kill(1), execve(2V), fork(2V), getitimer(2), getrlimit(2), kill(2V), ptrace(2), read(2V), sigblock(2), sigpause(2V), sigsetmask(2), sigstack(2), sigvec(2), vfork(2), wait(2V), write(2V), setjmp(3V), termio(4)

#### NOTES

The handler routine can be declared:

```
void handler(sig, code, scp, addr)
int sig, code;
struct sigcontext *scp;
char *addr;
```

Here *sig* is the signal number; *code* is a parameter of certain signals that provides additional detail; *scp* is a pointer to the sigcontext structure (defined in <signal.h>), used to restore the context from before the signal; and *addr* is additional address information. See sigvec(2) for more details.

**NAME**

sigsetops, sigaddset, sigdelset, sigfillset, sigemptyset, sigismember – manipulate signal sets

**SYNOPSIS**

```
#include <signal.h>

int sigaddset(set, signo)
sigset_t *set;
int signo;

int sigdelset(set, signo)
sigset_t *set;
int signo;

int sigfillset(set)
sigset_t *set;

int sigemptyset(set)
sigset_t *set;

int sigismember(set, signo)
sigset_t *set
int signo;
```

**DESCRIPTION**

The **sigsetops** primitives manipulate sets of signals. They operate on data objects addressable by the application. They do not operate on any set of signals known to the system, such as the set blocked from delivery to a process or the set pending for a process.

**sigaddset()** and **sigdelset()** respectively add and delete the individual signal specified by the value of **signo** from the signal set pointed to by **set**.

**sigemptyset()** initializes the signal set pointed to by **set** such that all signals defined in this standard are excluded.

**sigfillset()** initializes the signal set pointed to by **set** such that all signals defined in this standard are included.

Applications shall call either **sigemptyset()** or **sigfillset()** at least once for each object of type **sigset\_t** prior to any other use of that object. If such an object is not initialized in this way, but is nonetheless supplied as an argument to any of **sigaddset()**, **sigdelset()**, **sigismember()**, **sigaction()**, **sigprocmask()**, **sigpending()**, or **sigsuspend()** the results are undefined.

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

**RETURN VALUES**

**sigismember()** returns:

- 1 if the specified signal is a member of **set**.
- 0 if the specified signal is not a member of **set**.
- 1 if an error is detected, and sets **errno** to indicate the error.

The other functions return:

- 0 on success.
- 1 on failure and set **errno** to indicate the error.

**ERRORS**

For each of the following conditions, if the condition is detected, **sigaddset()**, **sigdelset()**, and **sigismember()** set **errno** to:

**EINVAL** **signo** is an invalid or unsupported signal number.

**SEE ALSO**

**sigaction(3V), sigpending(2V), sigprocmask(2V)**

**NAME**

sleep – suspend execution for interval

**SYNOPSIS**

**int sleep(seconds)**  
**unsigned seconds;**

**SYSTEM V SYNOPSIS**

**unsigned sleep(seconds)**  
**unsigned seconds;**

**DESCRIPTION**

**sleep()** suspends the current process from execution for the number of seconds specified by the argument. The actual suspension time may be an arbitrary amount longer because of other activity in the system.

**sleep()** is implemented by setting an interval timer and pausing until it expires. The previous state of this timer is saved and restored. If the sleep time exceeds the time to the expiration of the previous value of the timer, the process sleeps only until the timer would have expired, and the signal which occurs with the expiration of the timer is sent one second later.

**SYSTEM V DESCRIPTION**

**sleep()** suspends the current process from execution until either the number of real time seconds specified by *seconds* have elapsed or a signal is delivered to the calling process and its action is to invoke a signal-catching function or to terminate the process. The suspension time may be an arbitrary amount longer than requested because of other activity in the system. The value returned by **sleep()** will be the “unslept” amount (the requested time minus the time actually slept) in case the caller had an alarm set to go off earlier than the end of the requested **sleep()** time, or premature arousal due to another caught signal.

**RETURN VALUES**

**sleep()** returns no useful value.

**SYSTEM V RETURN VALUES**

If **sleep()** returns because the requested time has elapsed, it returns 0. If **sleep()** returns due to the delivery of a signal, it returns the “unslept” amount in seconds.

**SEE ALSO**

**getitimer(2)**, **sigpause(2V)**, **usleep(3)**

**NOTES**

SIGALRM should *not* be blocked or ignored during a call to **sleep()**. Only a prior call to **alarm(3V)** should generate SIGALRM for the calling process during a call to **sleep()**. A signal-catching function should *not* interrupt a call to **sleep()** to call **siglongjmp()** or **longjmp()** to restore an environment saved prior to the **sleep()** call.

**WARNINGS**

**sleep()** is slightly incompatible with **alarm(3V)**. Programs that do not execute for at least one second of clock time between successive calls to **sleep()** indefinitely delay the alarm signal. Use System V **sleep()**. Each **sleep(3V)** call postpones the alarm signal that would have been sent during the requested sleep period to occur one second later.



**NAME**

**sputl, sgetl** – access long integer data in a machine-independent fashion

**SYNOPSIS**

**void sputl(value, buffer)**

**long value;**

**char \*buffer;**

**long sgetl(buffer)**

**char \*buffer;**

**DESCRIPTION**

**sputl()** takes the four bytes of the long integer **values** and places them in memory starting at the address pointed to by *buffer*. The ordering of the bytes is the same across all machines.

**sgetl()** retrieves the four bytes in memory starting at the address pointed to by *buffer* and returns the long integer value in the byte ordering of the host machine.

The combination of **sputl()** and **sgetl()** provides a machine-independent way of storing long numeric data in a file in binary form without conversion to characters.

**NAME**

ssignal, gsignal – software signals

**SYNOPSIS**

```
#include <signal.h>
```

```
int (*ssignal (sig, action))()
```

```
int sig, (*action)();
```

```
int gsignal (sig)
```

```
int sig;
```

**DESCRIPTION**

**ssignal()** and **ssignal()** implement a software facility similar to **signal(3V)**.

Software signals made available to users are associated with integers in the inclusive range 1 through 15. A call to **ssignal()** associates a procedure, *action*, with the software signal *sig*; the software signal, *sig*, is raised by a call to **ssignal()**. Raising a software signal causes the action established for that signal to be *taken*.

The first argument to **ssignal()** is a number identifying the type of signal for which an action is to be established. The second argument defines the action; it is either the name of a (user-defined) *action function* or one of the manifest constants **SIG\_DFL** (default) or **SIG\_IGN** (ignore). **ssignal()** returns the action previously established for that signal type; if no action has been established or the signal number is illegal, **ssignal()** returns **SIG\_DFL**.

**ssignal()** raises the signal identified by its argument, *sig*:

If an action function has been established for *sig*, then that action is reset to **SIG\_DFL** and the action function is entered with argument *sig*. **ssignal()** returns the value returned to it by the action function.

If the action for *sig* is **SIG\_IGN**, **ssignal()** returns the value 1 and takes no other action.

If the action for *sig* is **SIG\_DFL**, **ssignal()** returns the value 0 and takes no other action.

If *sig* has an illegal value or no action was ever specified for *sig*, **ssignal()** returns the value 0 and takes no other action.

**SEE ALSO**

**signal(3V)**

**NAME**

stdio – standard buffered input/output package

**SYNOPSIS**

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

```
FILE *stdin;
```

```
FILE *stdout;
```

```
FILE *stderr;
```

**DESCRIPTION**

The functions described in section 3S constitute a user-level I/O buffering scheme. The in-line macros `getc(3V)` and `putc(3S)` handle characters quickly. The macros `getchar()` (see `getc(3V)`) and `putchar()` (see `putc(3S)`), and the higher level routines `fgetc()`, `getw()` (see `getc(3V)`), `gets(3S)`, `fgets()` (see `gets(3S)`), `scanf(3V)`, `fscanf()` (see `scanf(3V)`), `fread(3S)`, `fputc()`, `putw()` (see `putc(3S)`), `puts(3S)`, `fputs()` (see `puts(3S)`), `printf(3V)`, `fprintf()` (see `printf(3V)`), `fwrite()` (see `fread(3S)`) all use or act as if they use `getc()` and `putc()`. They can be freely intermixed.

A file with associated buffering is called a *stream*, and is declared to be a pointer to a defined type **FILE**. `fopen(3V)` creates certain descriptive data for a stream and returns a pointer to designate the stream in all further transactions. Normally, there are three open streams with constant pointers declared in the `<stdio.h>` include file and associated with the standard open files:

```
stdin      standard input file
stdout     standard output file
stderr     standard error file
```

A constant `NULL (0)` designates a nonexistent pointer.

An integer constant `EOF (-1)` is returned upon EOF or error by most integer functions that deal with streams (see the individual descriptions for details).

Any module that uses this package must include the header file of pertinent macro definitions, as follows:

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

The functions and constants mentioned in sections labeled 3S of this manual are declared in that header file and need no further declaration. The constants and the following ‘functions’ are implemented as macros; redeclaration of these names is perilous: `getc()`, `getchar()`, `putc()`, `putchar()`, `feof()`, `ferror()`, `fileno()`, and `clearerr()`.

Output streams, with the exception of the standard error stream `stderr`, are by default buffered if the output refers to a file and line-buffered if the output refers to a terminal. The standard error output stream `stderr` is by default unbuffered, but use of `fopen()` will cause it to become buffered or line-buffered. When an output stream is unbuffered, information is written to the destination file or terminal as soon as it is output to the stream; when it is buffered, many characters are saved up and written as a block. When it is line-buffered, each line of output is written to the destination file or terminal as soon as the line is completed (that is, as soon as a `NEWLINE` character is output or, if the output stream is `stdout` or `stderr`, as soon as input is read from `stdin`). `setbuf(3V)`, `setbuffer()`, `setlinebuf()`, or `setvbuf()` (see `setbuf(3V)`) can be used to change the stream’s buffering strategy.

**SYSTEM V DESCRIPTION**

When an output stream is line-buffered, each line of output is written to the destination file or terminal as soon as the line is completed (that is, as soon as a `NEWLINE` character is output or as soon as input is read from a line-buffered stream).

Output saved up on *all* line-buffered streams is written when input is read from *any* line-buffered stream. Input read from a stream that is not line-buffered does not flush output on line-buffered streams.

**RETURN VALUES**

The value EOF is returned uniformly to indicate that a FILE pointer has not been initialized with fopen(), input (output) has been attempted on an output (input) stream, or a FILE pointer designates corrupt or otherwise unintelligible FILE data.

**SEE ALSO**

open(2V), close(2V), lseek(2V), pipe(2V), read(2V), write(2V), ctermid(3V), cuserid(3V), fclose(3V), perror(3V), fopen(3V), fread(3S), fseek(3S), getc(3V), gets(3S), popen(3S), printf(3V), putc(3S), puts(3S), scanf(3V), setbuf(3V), system(3), tmpfile(3S), tmpnam(3S), ungetc(3S)

**NOTES**

The line buffering of output to terminals is almost always transparent, but may cause confusion or malfunctioning of programs which use standard I/O routines but use read(2V) to read from the standard input, as calls to read() do not cause output to line-buffered streams to be flushed.

In cases where a large amount of computation is done after printing part of a line on an output terminal, it is necessary to call fflush() (see fclose(3V)) on the standard output before performing the computation so that the output will appear.

**BUGS**

The standard buffered functions do not interact well with certain other library and system functions, especially vfork(2).

**NAME**

**strcoll, strxfrm** – compare or transform strings using collating information

**SYNOPSIS**

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

```
int strcoll(s1, s2)
```

```
char *s1;
```

```
char *s2;
```

```
size_t strxfrm(s1, s2, n)
```

```
char *s1;
```

```
char *s2;
```

```
size_t n;
```

**DESCRIPTION**

**strcoll()** compares the string pointed to by *s1* to the string pointed to by *s2*. These strings are interpreted as appropriate to the **LC\_COLLATE** category of the current locale.

**strxfrm()** transforms the string pointed to by *s2* and places the resulting string into the array pointed to by *s1*. The transformation is such that if **string()** is applied to two transformed strings, it returns a value greater than, equal to, or less than zero, corresponding to the result of the **strcoll()** function applied to the same two original strings. No more than *n* characters are placed into the resulting array pointed to by *s1*, including the terminating null character. If *n* is zero, *s1* is permitted to be a null pointer. If copying takes place between objects that overlap, the behavior is undefined.

**RETURN VALUES**

On success, **strcoll()** returns an integer greater than, equal to or less than zero, respectively, if the string pointed to by *s1* is greater than, equal to or less than the string pointed to by *s2* when both are interpreted as appropriate to the current locale. On failure, **strcoll()** sets **errno** to indicate the error, but returns no special value.

**strxfrm()** returns the length of the transformed string, not including the terminating null character. If the value returned is *n* or more, the contents of the array pointed to by *s1* are indeterminate. On failure, **strxfrm()** returns **(size\_t)-1**, and sets **errno** to indicate the error.

**ERRORS**

**EINVAL**            *s1* or *s2* contain characters outside the domain of the collating sequence.

**SEE ALSO**

**string(3)**

## NAME

strcat, strncat, strdup, strcmp, strncmp, strcasecmp, strncasecmp, strcpy, strncpy, strlen, strchr, strrchr, strpbrk, strspn, strcspn, strstr, strtok, index, rindex – string operations

## SYNOPSIS

```
#include <string.h>

char *strcat(s1, s2)
char *s1, *s2;

char *strncat(s1, s2, n)
char *s1, *s2;
int n;

char *strdup(s1)
char *s1;

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

int strncmp(s1, s2, n)
char *s1, *s2;
int n;

int strcasecmp(s1, s2) char *s1, *s2;

int strncasecmp(s1, s2, n)
char *s1, *s2;
int n;

char *strcpy(s1, s2)
char *s1, *s2;

char *strncpy(s1, s2, n)
char *s1, *s2;
int n;

int strlen(s)
char *s;

char *strchr(s, c)
char *s;
int c;

char *strrchr(s, c)
char *s;
int c;

char *strpbrk(s1, s2)
char *s1, *s2;

int strspn(s1, s2)
char *s1, *s2;

int strcspn(s1, s2)
char *s1, *s2;

char *strstr(s1, s2)
char *s1, *s2;

char *strtok(s1, s2)
char *s1, *s2;
```

```
#include <strings.h>
char *index(s, c)
char *s, c;
char *rindex(s, c)
char *s, c;
```

## DESCRIPTION

These functions operate on null-terminated strings. They do not check for overflow of any receiving string. **strcat()** appends a copy of string *s2* to the end of string *s1*. **strncat()** appends at most *n* characters. Each returns a pointer to the null-terminated result.

**strcmp()** compares its arguments and returns an integer greater than, equal to, or less than 0, according as *s1* is lexicographically greater than, equal to, or less than *s2*. **strncmp()** makes the same comparison but compares at most *n* characters. Two additional routines **strcasecmp()** and **strncasecmp()** compare the strings and ignore differences in case. These routines assume the ASCII character set when equating lower and upper case characters.

**strdup()** returns a pointer to a new string which is a duplicate of the string pointed to by *s1*. The space for the new string is obtained using **malloc(3V)**. If the new string cannot be created, a NULL pointer is returned.

**strcpy()** copies string *s2* to *s1* until the null character has been copied. **strncpy()** copies string *s2* to *s1* until either the null character has been copied or *n* characters have been copied. If the length of *s2* is less than *n*, **strncpy()** pads *s1* with null characters. If the length of *s2* is *n* or greater, *s1* will not be null-terminated. Both functions return *s1*.

**strlen()** returns the number of characters in *s*, not including the null-terminating character.

**strchr()** (**strrchr()**) returns a pointer to the first (last) occurrence of character *c* in string *s*, or a NULL pointer if *c* does not occur in the string. The null character terminating a string is considered to be part of the string.

**index()** (**rindex()**) returns a pointer to the first (last) occurrence of character *c* in string *s*, or a NULL pointer if *c* does not occur in the string. These functions are identical to **strchr()** (**strrchr()**) and merely have different names.

**strpbrk()** returns a pointer to the first occurrence in string *s1* of any character from string *s2*, or a NULL pointer if no character from *s2* exists in *s1*.

**strspn()** (**strcspn()**) returns the length of the initial segment of string *s1* which consists entirely of characters from (not from) string *s2*.

**strstr()** returns a pointer to the first occurrence of the pattern string *s2* in *s1*. For example, if *s1* is "string thing" and *s2* is "ing", **strstr()** returns "ing thing". If *s2* does not occur in *s1*, **strstr()** returns NULL.

**strtok()** considers the string *s1* to consist of a sequence of zero or more text tokens separated by spans of one or more characters from the separator string *s2*. The first call (with pointer *s1* specified) returns a pointer to the first character of the first token, and will have written a null character into *s1* immediately following the returned token. The function keeps track of its position in the string between separate calls, so that subsequent calls (which must be made with the first argument a NULL pointer) will work through the string *s1* immediately following that token. In this way subsequent calls will work through the string *s1* until no tokens remain. The separator string *s2* may be different from call to call. When no token remains in *s1*, a NULL pointer is returned.

**NOTES**

For user convenience, all these functions, except for **index()** and **rindex()**, are declared in the optional **<string.h>** header file. All these functions, including **index()** and **rindex()** but excluding **strchr()**, **strchr()**, **strpbrk()**, **strspn()**, **strcspn()**, and **strtok()** are declared in the optional **<strings.h>** include file; these headers are set this way for backward compatibility.

**SEE ALSO**

**malloc(3V)**, **bstring(3)**

**WARNINGS**

**strcmp()** and **strncmp()** use native character comparison, which is signed on the Sun, but may be unsigned on other machines. Thus the sign of the value returned when one of the characters has its high-order bit set is implementation-dependent.

**strcasecmp()** and **strncasecmp()** use native character comparison as above and assume the *ASCII* character set.

On the Sun processor, as well as on many other machines, you can *not* use a NULL pointer to indicate a null string. A NULL pointer is an error and results in an abort of the program. If you wish to indicate a null string, you must have a pointer that points to an explicit null string. On some implementations of the C language on some machines, a NULL pointer, if dereferenced, would yield a null string; this highly non-portable trick was used in some programs. Programmers using a NULL pointer to represent an empty string should be aware of this portability issue; even on machines where dereferencing a NULL pointer does not cause an abort of the program, it does not necessarily yield a null string.

Character movement is performed differently in different implementations. Thus overlapping moves may yield surprises.



## NAME

string\_to\_decimal, file\_to\_decimal, func\_to\_decimal – parse characters into decimal record

## SYNOPSIS

```
#include <floatingpoint.h>
```

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

```
void string_to_decimal(pc,nmax,fortran_conventions,pd,pform,pechar)
```

```
char **pc;
```

```
int nmax;
```

```
int fortran_conventions;
```

```
decimal_record *pd;
```

```
enum decimal_string_form *pform;
```

```
char **pechar;
```

```
void file_to_decimal(pc,nmax,fortran_conventions,pd,pform,pechar,pf,pnread)
```

```
char **pc;
```

```
int nmax;
```

```
int fortran_conventions;
```

```
decimal_record *pd;
```

```
enum decimal_string_form *pform;
```

```
char **pechar;
```

```
FILE *pf;
```

```
int *pnread;
```

```
void func_to_decimal(pc,nmax,fortran_conventions,pd,pform,pechar,pget,pnread,punget)
```

```
char **pc;
```

```
int nmax;
```

```
int fortran_conventions;
```

```
decimal_record *pd;
```

```
enum decimal_string_form *pform;
```

```
char **pechar;
```

```
int (*pget)();
```

```
int *pnread;
```

```
int (*punget)();
```

## DESCRIPTION

The `char_to_decimal()` functions parse a numeric token from at most *nmax* characters in a string *\*\*pc* or file *\*pf* or function *(\*pget)()* into a decimal record *\*pd*, classifying the form of the string in *\*pform* and *\*pechar*. The accepted syntax is intended to be sufficiently flexible to accommodate many languages:

*whitespace value*

or

*whitespace sign value*

where *whitespace* is any number of characters defined by *isspace* in `<ctype.h>`, *sign* is either of `[+-]`, and *value* can be *number*, *nan*, or *inf*. *inf* can be `INF` (*inf\_form*) or `INFINITY` (*infinity\_form*) without regard to case. *nan* can be `NAN` (*nan\_form*) or `NAN(nstring)` (*nanstring\_form*) without regard to case; *nstring* is any string of characters not containing `'`' or the null character; *nstring* is copied to *pd*→*ds* and, currently, not used subsequently. *number* consists of

*significant*

or

*significant efield*

where *significant* must contain one or more digits and may contain one point; possible forms are

<i>digits</i>	( <i>int_form</i> )
<i>digits.</i>	( <i>intdot_form</i> )
<i>.digits</i>	( <i>dotfrac_form</i> )
<i>digits.digits</i>	( <i>intdotfrac_form</i> )

*efield* consists of

*echar digits*

or

*echar sign digits*

where *echar* is one of [Ee], and *digits* contains one or more digits.

When *fortran\_conventions* is nonzero, additional input forms are accepted according to various Fortran conventions:

- 0 no Fortran conventions
- 1 Fortran list-directed input conventions
- 2 Fortran formatted input conventions, ignore blanks (BN)
- 3 Fortran formatted input conventions, blanks are zeros (BZ)

When *fortran\_conventions* is nonzero, *echar* may also be one of [Dd], and *efield* may also have the form

*sign digits.*

When *fortran\_conventions*  $\geq$  2, blanks may appear in the *digits* strings for the integer, fraction, and exponent fields and may appear between *echar* and the exponent sign and after the infinity and NaN forms. If *fortran\_conventions*  $=$  2, the blanks are ignored. When *fortran\_conventions*  $=$  3, the blanks that appear in *digits* strings are interpreted as zeros, and other blanks are ignored.

When *fortran\_conventions* is zero, the current locale's decimal point character is used as the decimal point; when *fortran\_conventions* is nonzero, the period is used as the decimal point.

The form of the accepted decimal string is placed in *\*perform*. If an *efield* is recognized, *\*pechar* is set to point to the *echar*.

On input, *\*pc* points to the beginning of a character string buffer of length  $\geq$  *nmax*. On output, *\*pc* points to a character in that buffer, one past the last accepted character. **string\_to\_decimal()** gets its characters from the buffer; **file\_to\_decimal()** gets its characters from *\*pf* and records them in the buffer, and places a null after the last character read. **func\_to\_decimal()** gets its characters from an int function (*\*pget*)().

The scan continues until no more characters could possibly fit the acceptable syntax or until *nmax* characters have been scanned. If the *nmax* limit is not reached then at least one extra character will usually be scanned that is not part of the accepted syntax. **file\_to\_decimal()** and **func\_to\_decimal()** set *\*pnread* to the number of characters read from the file; if greater than *nmax*, some characters were lost. If no characters were lost, **file\_to\_decimal()** and **func\_to\_decimal()** attempt to push back, with **ungetc(3S)** or (**\*punget**)(), as many as possible of the excess characters read, adjusting *\*pnread* accordingly. If all **ungetc** calls are successful, then *\*\*pc* will be a null character. No push back will be attempted if (**\*punget**)() is NULL.

Typical declarations for *\*pget()* and *\*punget()* are:

```
int xget()
{ ... }
int (*pget)() = xget;
int xunget(c)
char c ;
{ ... }
int (*punget)() = xunget;
```

If no valid number was detected, *pd->fpclass* is set to **fp\_signaling**, *\*pc* is unchanged, and *\*pform* is set to **invalid\_form**.

**atof()** and **strtod(3)** use **string\_to\_decimal()**. **scanf(3V)** uses **file\_to\_decimal()**.

**SEE ALSO**

**ctype(3V)**, **localeconv(3)**, **scanf(3V)**, **setlocale(3V)**, **strtod(3)**, **ungetc(3S)**

**NAME**

**strtod, atof** – convert string to double-precision number

**SYNOPSIS**

```
double strtod(str, ptr)
```

```
char *str, **ptr;
```

```
double atof(str)
```

```
char *str;
```

**DESCRIPTION**

**strtod()** returns as a double-precision floating-point number the value represented by the character string pointed to by *str*. The string is scanned up to the first unrecognized character, using **string\_to\_decimal(3)**, with *fortran\_conventions* set to 0.

If the value of *ptr* is not (char \*\*)NULL, a pointer to the character terminating the scan is returned in the location pointed to by *ptr*. If no number can be formed, *\*ptr* is set to *str*, and for historical compatibility, 0.0 is returned, although a NaN would better match the IEEE Floating-Point Standard's intent.

The radix character is defined by the program's locale (category LC\_NUMERIC). In the "C" locale, or in a locale where the radix character is not defined, the radix character defaults to a period '.'.

**atof(str)** is equivalent to **strtod(str, (char \*\*)NULL)**. Thus, when **atof(str)** returns 0.0 there is no way to determine whether *str* contained a valid numerical string representing 0.0 or an invalid numerical string.

**SEE ALSO**

**scanf(3V)**, **string\_to\_decimal(3)**

**DIAGNOSTICS**

Exponent overflow and underflow produce the results specified by the IEEE Standard. In addition, **errno** is set to ERANGE.

**NAME**

`strtol`, `atol`, `atoi` – convert string to integer

**SYNOPSIS**

**long** `strtol`(*str*, *ptr*, *base*)

**char \****str*, **\*\****ptr*;

**int** *base*;

**long** `atol`(*str*)

**char \****str*;

**int** `atoi`(*str*)

**char \****str*;

**DESCRIPTION**

`strtol()` returns as a long integer the value represented by the character string pointed to by *str*. The string is scanned up to the first character inconsistent with the base. Leading “white-space” characters (as defined by `isspace()` in `ctype(3V)`) are ignored.

If the value of *ptr* is not `(char **)NULL`, a pointer to the character terminating the scan is returned in the location pointed to by *ptr*. If no integer can be formed, that location is set to *str*, and zero is returned.

If *base* is positive (and not greater than 36), it is used as the base for conversion. After an optional leading sign, leading zeros are ignored, and “0x” or “0X” is ignored if *base* is 16.

If *base* is zero, the string itself determines the base thusly: after an optional leading sign a leading zero indicates octal conversion, and a leading “0x” or “0X” hexadecimal conversion. Otherwise, decimal conversion is used.

Truncation from long to int can, of course, take place upon assignment or by an explicit cast.

`atol(str)` is equivalent to `strtol(str, (char **)NULL, 10)`.

`atoi(str)` is equivalent to `(int) strtol(str, (char **)NULL, 10)`.

**SEE ALSO**

`ctype(3V)`, `scanf(3V)`, `strtod(3)`

**BUGS**

Overflow conditions are ignored.

**NAME**

stty, gtty – set and get terminal state

**SYNOPSIS**

```
#include <sgtty.h>
stty(fd, buf)
int fd;
struct sgttyb *buf;
gtty(fd, buf)
int fd;
struct sgttyb *buf;
```

**DESCRIPTION**

Note: this interface is obsoleted by `ioctl(2)`.

`stty()` sets the state of the terminal associated with *fd*. `stty()` retrieves the state of the terminal associated with *fd*. To set the state of a terminal the call must have write permission.

The `stty()` call is actually

```
ioctl(fd, TIOCSETP, buf)
```

while the `gtty()` call is

```
ioctl(fd, TIOCGETP, buf)
```

See `ioctl(2)` and `ttcompat(4M)` for an explanation.

**DIAGNOSTICS**

If the call is successful 0 is returned, otherwise -1 is returned and the global variable `errno` contains the reason for the failure.

**SEE ALSO**

`ioctl(2)`, `ttcompat(4M)`

**NAME**

swab – swap bytes

**SYNOPSIS**

```
void  
swab(from, to, nbytes)  
char *from, *to;
```

**DESCRIPTION**

**swab()** copies *nbytes* bytes pointed to by *from* to the position pointed to by *to*, exchanging adjacent even and odd bytes. It is useful for carrying binary data between high-ender machines (IBM 360's, MC68000's, etc) and low-end machines (such as Sun386i systems).

*nbytes* should be even and positive. If *nbytes* is odd and positive, **swab()** uses *nbytes* – 1 instead. If *nbytes* is negative, **swab()** does nothing.

The *from* and *to* addresses should not overlap in portable programs.

## NAME

syslog, openlog, closelog, setlogmask – control system log

## SYNOPSIS

```
#include <syslog.h>

openlog(ident, logopt, facility)
char *ident;

syslog(priority, message, parameters ... )
char *message;

closelog()

setlogmask(maskpri)
```

## DESCRIPTION

`syslog()` passes *message* to `syslogd(8)`, which logs it in an appropriate system log, writes it to the system console, forwards it to a list of users, or forwards it to the `syslogd` on another host over the network. The message is tagged with a priority of *priority*. The message looks like a `printf(3V)` string except that `%m` is replaced by the current error message (collected from `errno`). A trailing NEWLINE is added if needed.

Priorities are encoded as a *facility* and a *level*. The facility describes the part of the system generating the message. The level is selected from an ordered list:

<code>LOG_EMERG</code>	A panic condition. This is normally broadcast to all users.
<code>LOG_ALERT</code>	A condition that should be corrected immediately, such as a corrupted system database.
<code>LOG_CRIT</code>	Critical conditions, such as hard device errors.
<code>LOG_ERR</code>	Errors.
<code>LOG_WARNING</code>	Warning messages.
<code>LOG_NOTICE</code>	Conditions that are not error conditions, but that may require special handling.
<code>LOG_INFO</code>	Informational messages.
<code>LOG_DEBUG</code>	Messages that contain information normally of use only when debugging a program.

If special processing is needed, `openlog()` can be called to initialize the log file. The parameter *ident* is a string that is prepended to every message. *logopt* is a bit field indicating logging options. Current values for *logopt* are:

<code>LOG_PID</code>	Log the process ID with each message. This is useful for identifying specific daemon processes (for daemons that fork).
<code>LOG_CONS</code>	Write messages to the system console if they cannot be sent to <code>syslogd</code> . This option is safe to use in daemon processes that have no controlling terminal, since <code>syslog()</code> forks before opening the console.
<code>LOG_NDELAY</code>	Open the connection to <code>syslogd</code> immediately. Normally the open is delayed until the first message is logged. This is useful for programs that need to manage the order in which file descriptors are allocated.
<code>LOG_NOWAIT</code>	Do not wait for child processes that have been forked to log messages onto the console. This option should be used by processes that enable notification of child termination using <code>SIGCHLD</code> , since <code>syslog()</code> may otherwise block waiting for a child whose exit status has already been collected.



The *facility* parameter encodes a default facility to be assigned to all messages that do not have an explicit facility already encoded:

<b>LOG_KERN</b>	Messages generated by the kernel. These cannot be generated by any user processes.
<b>LOG_USER</b>	Messages generated by random user processes. This is the default facility identifier if none is specified.
<b>LOG_MAIL</b>	The mail system.
<b>LOG_DAEMON</b>	System daemons, such as <b>ftpd(8C)</b> , <b>routed(8C)</b> , etc.
<b>LOG_AUTH</b>	The authorization system: <b>login(1)</b> , <b>su(1V)</b> , <b>getty(8)</b> , etc.
<b>LOG_LPR</b>	The line printer spooling system: <b>lpr(1)</b> , <b>lpc(8)</b> , <b>lpd(8)</b> , etc.
<b>LOG_NEWS</b>	Reserved for the USENET network news system.
<b>LOG_UUCP</b>	Reserved for the UUCP system; it does not currently use <b>syslog</b> .
<b>LOG_CRON</b>	The <b>cron/at</b> facility; <b>crontab(1)</b> , <b>at(1)</b> , <b>cron(8)</b> , etc.
<b>LOG_LOCAL0-7</b>	Reserved for local use.

**closelog()** can be used to close the log file.

**setlogmask()** sets the log priority mask to *maskpri* and returns the previous mask. Calls to **syslog()** with a priority not set in *maskpri* are rejected. The mask for an individual priority *pri* is calculated by the macro **LOG\_MASK(pri)**; the mask for all priorities up to and including *toppri* is given by the macro **LOG\_UPTO(toppri)**. The default allows all priorities to be logged.

#### EXAMPLES

This call logs a message at priority **LOG\_ALERT**:

```
syslog(LOG_ALERT, "who: internal error 23");
```

The FTP daemon **ftpd** would make this call to **openlog()** to indicate that all messages it logs should have an identifying string of **ftpd**, should be treated by **syslogd** as other messages from system daemons are, should include the process ID of the process logging the message:

```
openlog("ftpd", LOG_PID, LOG_DAEMON);
```

Then it would make the following call to **setlogmask()** to indicate that messages at priorities from **LOG\_EMERG** through **LOG\_ERR** should be logged, but that no messages at any other priority should be logged:

```
setlogmask(LOG_UPTO(LOG_ERR));
```

Then, to log a message at priority **LOG\_INFO**, it would make the following call to **syslog**:

```
syslog(LOG_INFO, "Connection from host %d", CallingHost);
```

A locally-written utility could use the following call to **syslog()** to log a message at priority **LOG\_INFO** to be treated by **syslogd** as other messages to the facility **LOG\_LOCAL2** are:

```
syslog(LOG_INFO|LOG_LOCAL2, "error: %m");
```

#### SEE ALSO

**at(1)**, **crontab(1)**, **logger(1)**, **login(1)**, **lpr(1)**, **su(1V)**, **printf(3V)**, **syslog.conf(5)**, **cron(8)**, **ftpd(8C)**, **getty(8)**, **lpc(8)**, **lpd(8)**, **routed(8C)**, **syslogd(8)**

**NAME**

system – issue a shell command

**SYNOPSIS**

```
system(string)
char *string;
```

**DESCRIPTION**

**system()** gives the *string* to **sh(1)** as input, just as if the string had been typed as a command from a terminal. The current process performs a **wait(2V)** system call, and waits until the shell terminates. **system()** then returns the exit status returned by **wait(2V)**. Unless the shell was interrupted by a signal, its termination status is contained in the 8 bits higher up from the low-order 8 bits of the value returned by **wait()**.

**SEE ALSO**

**sh(1)**, **execve(2V)**, **wait(2V)**, **popen(3S)**

**DIAGNOSTICS**

Exit status 127 (may be displayed as "32512") indicates the shell could not be executed.

## NAME

`t_accept` – accept a connect request

## SYNOPSIS

```
#include <tiuser.h>

int t_accept(fd, resfd, call)
int fd;
int resfd;
struct t_call *call;
```

## DESCRIPTION

`t_accept()` is issued by a transport user to accept a connect request. *fd* identifies the local transport endpoint where the connect indication arrived, *resfd* specifies the local transport endpoint where the connection is to be established, and *call* contains information required by the transport provider to complete the connection. *call* points to a `t_call` structure which contains the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
int sequence;
```

The *netbuf* structure contains the following members:

```
unsigned int maxlen;
unsigned int len;
char *buf;
```

*buf* points to a user input and/or output buffer. *len* generally specifies the number of bytes contained in the buffer. If the structure is used for both input and output, the transport function will replace the user value of *len* on return. *maxlen* generally has significance only when *buf* is used to receive output from the transport function. In this case, it specifies the physical size of the buffer, and the maximum value of *len* that can be set by the function. If *maxlen* is not large enough to hold the returned information, a TBUFOVFLW error will generally result. However, certain functions may return part of the data and not generate an error. In *call*, *addr* is the address of the caller, *opt* indicates any protocol-specific parameters associated with the connection, *udata* points to any user data to be returned to the caller, and *sequence* is the value returned by `t_listen(3N)` that uniquely associates the response with a previously received connect indication.

A transport user may accept a connection on either the same, or on a different, local transport endpoint than the one on which the connect indication arrived. If the same endpoint is specified (*resfd* = *fd*), the connection can be accepted unless the following condition is true: The user has received other indications on that endpoint but has not responded to them (with `t_accept()` or `t_snddis(3N)`). For this condition, `t_accept()` will fail and set `t_errno` to TBADF.

If a different transport endpoint is specified (*resfd* != *fd*), the endpoint must be bound to a protocol address and must be in the T\_IDLE state (see `t_getstate(3N)`) before the `t_accept()` is issued.

For both types of endpoints, `t_accept()` will fail and set `t_errno` to TLOOK if there are indications (such as a connect or disconnect) waiting to be received on that endpoint.

The values of parameters specified by *opt* and the syntax of those values are protocol specific. The *udata* field enables the called transport user to send user data to the caller and the amount of user data must not exceed the limits supported by the transport provider as returned by `t_open(3N)` or `t_getinfo(3N)`. If the *len* field of *udata* is zero, no data will be sent to the caller.

## RETURN VALUES

`t_accept()` returns:

```
0      on success.
-1     on failure and sets t_errno to indicate the error.
```

**ERRORS**

TACCES	The user does not have permission to accept a connection on the responding transport endpoint. The user does not have permission to use the specified options.
TBADDATA	The amount of user data specified was not within the bounds allowed by the transport provider.
TBADF	The specified file descriptor does not refer to a transport endpoint. The user is illegally accepting a connection on the same transport endpoint on which the connect indication arrived.
TBADOPT	The specified options were in an incorrect format or contained illegal information.
TBADSEQ	An invalid sequence number was specified.
TLOOK	An asynchronous event has occurred on the transport endpoint referenced by <i>fd</i> and requires immediate attention.
TNOTSUPPORT	This function is not supported by the underlying transport provider.
TOUTSTATE	The function was issued in the wrong sequence on the transport endpoint referenced by <i>fd</i> . The transport endpoint referred to by <i>resfd</i> is not in the T_IDLE state.
TSYSERR	The function failed due to a system error and set <code>errno</code> to indicate the error.

**SEE ALSO**

`intro(3)`, `t_connect(3N)`, `t_getstate(3N)`, `t_listen(3N)`, `t_open(3N)`, `t_rcvconnect(3N)`

*Network Programming*

## NAME

`t_alloc` – allocate a library structure

## SYNOPSIS

```
#include <tiuser.h>

char *t_alloc(fd, struct_type, fields)
int fd;
int struct_type;
int fields;
```

## DESCRIPTION

`t_alloc()` dynamically allocates memory for the various transport function argument structures as specified below. `t_alloc()` allocates memory for the specified structure and for buffers referenced by the structure.

The structure to allocate is specified by `struct_type`, and can be one of the following (each of these structures may be used as an argument to one or more transport functions):

<code>T_BIND</code>	<code>struct t_bind</code>
<code>T_CALL</code>	<code>struct t_call</code>
<code>T_OPTMGMT</code>	<code>struct t_optmgmt</code>
<code>T_DIS</code>	<code>struct t_discon</code>
<code>T_UNITDATA</code>	<code>struct t_unitdata</code>
<code>T_UDERROR</code>	<code>struct t_uderr</code>
<code>T_INFO</code>	<code>struct t_info</code>

Each of the above structures, except `T_INFO`, contains at least one field of type ‘`struct netbuf`’. The `maxlen`, `len`, and `buf` members of the `netbuf` structure are described in `t_accept(3N)`. For each field of this type, the user may specify that the buffer for that field should be allocated as well. The `fields` argument specifies this option, where the argument is the bitwise-OR of any of the following:

<code>T_ADDR</code>	The <code>addr</code> field of the <code>t_bind</code> , <code>t_call</code> , <code>t_unitdata</code> , or <code>t_uderr</code> structures.
<code>T_OPT</code>	The <code>opt</code> field of the <code>t_optmgmt</code> , <code>t_call</code> , <code>t_unitdata</code> , or <code>t_uderr</code> structures.
<code>T_UDATA</code>	The <code>udata</code> field of the <code>t_call</code> , <code>t_discon</code> , or <code>t_unitdata</code> structures.
<code>T_ALL</code>	All relevant fields of the given structure.

For each field specified in `fields`, `t_alloc()` allocates memory for the buffer associated with the field, and initializes the `buf` pointer and `maxlen` field accordingly. The length of the buffer allocated is based on the same size information returned to the user on `t_open(3N)` and `t_getinfo(3N)`. Thus, `fd` must refer to the transport endpoint through which the newly allocated structure is passed, so that the appropriate size information can be accessed. If the size value associated with any specified field is `-1` or `-2` (see `t_open(3N)` or `t_getinfo(3N)`), `t_alloc()` is unable to determine the size of the buffer to allocate and fails, setting `t_errno` to `TSYSERR` and `errno` to `EINVAL`. For any field not specified in `fields`, `buf` is set to `NULL` and `maxlen` is set to zero.

Use of `t_alloc()` to allocate structures helps ensure the compatibility of user programs with future releases of the transport interface.

## RETURN VALUES

On success, `t_alloc()` returns a pointer to the type of structure specified by `struct_type`. On failure, it returns `NULL` and sets `t_errno` to indicate the error.

## ERRORS

<code>TBADF</code>	The specified file descriptor does not refer to a transport endpoint.
<code>TSYSERR</code>	The function failed due to a system error and set <code>errno</code> to indicate the error.

**SEE ALSO**

**intro(3), t\_free(3N), t\_getinfo(3N), t\_open(3N)**

*Network Programming*

## NAME

**t\_bind** – bind an address to a transport endpoint

## SYNOPSIS

```
#include <tiuser.h>

int t_bind(fd, req, ret)
int fd;
struct t_bind *req;
struct t_bind *ret;
```

## DESCRIPTION

**t\_bind()** associates a protocol address with the transport endpoint specified by *fd* and activates that transport endpoint. In connection mode, the transport provider may begin accepting or requesting connections on the transport endpoint. In connectionless mode, the transport user may send or receive data units through the transport endpoint.

The *req* and *ret* arguments point to a **t\_bind()** structure containing the following members:

```
struct netbuf addr;
unsigned qlen;
```

The *maxlen*, *len*, and *buf* members of the *netbuf* structure are described in **t\_accept(3N)**. The *addr* field of the **t\_bind()** structure specifies a protocol address and the *qlen* field is used to indicate the maximum number of outstanding connect indications.

*req* is used to request that an address, represented by the *netbuf* structure, be bound to the given transport endpoint. *len* specifies the number of bytes in the address and *buf* points to the address buffer. *maxlen* has no meaning for the *req* argument. On return, *ret* contains the address that the transport provider actually bound to the transport endpoint; this may be different from the address specified by the user in *req*. In *ret*, the user specifies *maxlen* which is the maximum size of the address buffer and *buf* which points to the buffer where the address is to be placed. On return, *len* specifies the number of bytes in the bound address and *buf* points to the bound address. If *maxlen* is not large enough to hold the returned address, an error will result.

If the requested address is not available, or if no address is specified in *req* (the *len* field of *addr* in *req* is 0) the transport provider will assign an appropriate address to be bound, and will return that address in the *addr* field of *ret*. The user can compare the addresses in *req* and *ret* to determine whether the transport provider bound the transport endpoint to a different address than that requested.

*req* may be NULL if the user does not wish to specify an address to be bound. Here, the value of *qlen* is assumed to be 0, and the transport provider must assign an address to the transport endpoint. Similarly, *ret* may be NULL if the user does not care what address was bound by the transport provider and is not interested in the negotiated value of *qlen*. It is valid to set *req* and *ret* to NULL for the same call, in which case the transport provider chooses the address to bind to the transport endpoint and does not return that information to the user.

The *qlen* field has meaning only when initializing a connection-mode service. It specifies the number of outstanding connect indications the transport provider should support for the given transport endpoint. An outstanding connect indication is one that has been passed to the transport user by the transport provider. A value of *qlen* greater than 0 is only meaningful when issued by a passive transport user that expects other users to call it. The value of *qlen* will be negotiated by the transport provider and may be changed if the transport provider cannot support the specified number of outstanding connect indications. On return, the *qlen* field in *ret* will contain the negotiated value.

**t\_bind()** allows more than one transport endpoint to be bound to the same protocol address (however, the transport provider must support this capability also), but binding more than one protocol address to the same transport endpoint is not allowed. If a user binds more than one transport endpoint to the same protocol address, only one endpoint can be used to listen for connect indications associated with that protocol address. In other words, only one **t\_bind()** for a given protocol address may specify a value of *qlen* greater than 0. In this way, the transport provider can identify which transport endpoint

should be notified of an incoming connect indication. If a user attempts to bind a protocol address to a second transport endpoint with a value of *qlen* greater than 0, the transport provider will assign another address to be bound to that endpoint. If a user accepts a connection on the transport endpoint that is being used as the listening endpoint, the bound protocol address will be found to be busy for the duration of that connection. No other transport endpoints may be bound for listening while that initial listening endpoint is in the data transfer phase. This will prevent more than one transport endpoint bound to the same protocol address from accepting connect indications.

**RETURN VALUES**

**t\_bind()** returns:

- 0        on success.
- 1       on failure and sets **t\_errno** to indicate the error.

**ERRORS**

- |           |   |
|-----------|---|
| TACCES    | The user does not have permission to use the specified address.   |
| TBADADDR  | The specified protocol address was in an incorrect format or contained illegal information.   |
| TBADF     | The specified file descriptor does not refer to a transport endpoint.   |
| TBUFOVFLW | The number of bytes allowed for an incoming argument is not sufficient to store the value of that argument. The transport provider's state will change to <b>T_IDLE</b> and the information to be returned in <i>ret</i> will be discarded. |
| TNOADDR   | The transport provider could not allocate an address.   |
| TOUTSTATE | The function was issued in the wrong sequence.  |
| TSYSERR   | The function failed due to a system error and set <b>errno</b> to indicate the error.   |

**SEE ALSO**

**intro(3)**, **t\_open(3N)**, **t\_optmgmt(3N)**, **t\_unbind(3N)**

*Network Programming*



**NAME**

**t\_close** – close a transport endpoint

**SYNOPSIS**

```
#include <tiuser.h>
```

```
int t_close(fd)
```

```
int fd;
```

**DESCRIPTION**

**t\_close()** informs the transport provider that the user is finished with the transport endpoint specified by *fd*, and frees any local library resources associated with the endpoint. In addition, **t\_close()** closes the file associated with the transport endpoint.

**t\_close()** should be called from the T\_UNBND state (see **t\_getstate(3N)**). However, **t\_close()** does not check state information, so it may be called from any state to close a transport endpoint. If this occurs, the local library resources associated with the endpoint will be freed automatically. In addition, **close(2V)** will be issued for that file descriptor; the close will be abortive if no other process has that file open, and will break any transport connection that may be associated with that endpoint.

**RETURN VALUES**

**t\_close()** returns:

0           on success.

-1          on failure and sets **t\_errno** to indicate the error.

**ERRORS**

TBADF           The specified file descriptor does not refer to a transport endpoint.

**SEE ALSO**

**close(2V)**, **t\_getstate(3N)**, **t\_open(3N)**, **t\_unbind(3N)**

*Network Programming*

## NAME

`t_connect` – establish a connection with another transport user

## SYNOPSIS

```
#include <tiuser.h>

int t_connect(fd, sndcall, rcvcall)
int fd;
struct t_call *sndcall;
struct t_call *rcvcall;
```

## DESCRIPTION

`t_connect()` enables a transport user to request a connection to the specified destination transport user. *fd* identifies the local transport endpoint where communication will be established, while *sndcall* and *rcvcall* point to a `t_call()` structure which contains the following members:

```
    struct netbuf addr;
    struct netbuf opt;
    struct netbuf udata;
    int sequence;
```

*sndcall* specifies information needed by the transport provider to establish a connection and *rcvcall* specifies information that is associated with the newly established connection.

The *maxlen*, *len*, and *buf* members of the *netbuf* structure are described in `t_accept(3N)`. In *sndcall*, *addr* specifies the protocol address of the destination transport user, *opt* presents any protocol-specific information that might be needed by the transport provider, *udata* points to optional user data that may be passed to the destination transport user during connection establishment, and *sequence* has no meaning for this function.

On return in *rcvcall*, *addr* returns the protocol address associated with the responding transport endpoint, *opt* presents any protocol-specific information associated with the connection, *udata* points to optional user data that may be returned by the destination transport user during connection establishment, and *sequence* has no meaning for this function.

*opt* implies no structure on the options that may be passed to the transport provider. The transport provider is free to specify the structure of any options passed to it. These options are specific to the underlying protocol of the transport provider. The user may choose not to negotiate protocol options by setting the *len* field of *opt* to 0. In this case, the transport provider may use default options.

*udata* enables the caller to pass user data to the destination transport user and receive user data from the destination user during connection establishment. However, the amount of user data must not exceed the limits supported by the transport provider as returned by `t_open(3N)` or `t_getinfo(3N)`. If the *len* field of *udata* is 0 in *sndcall*, no data will be sent to the destination transport user.

On return, the *addr*, *opt*, and *udata* fields of *rcvcall* will be updated to reflect values associated with the connection. Thus, the *maxlen* field of each argument must be set before issuing this function to indicate the maximum size of the buffer for each. However, *rcvcall* may be NULL in which case no information is given to the user on return from `t_connect()`.

By default, `t_connect()` executes in synchronous mode, and will wait for the destination user's response before returning control to the local user. A successful return (a return value of 0) indicates that the requested connection has been established. However, if `T_NDELAY` is set (using `t_open()` or `fcntl()`), `t_connect()` executes in asynchronous mode. In this case, the call will not wait for the remote user's response, but will return control immediately to the local user and return -1 with `t_errno` set to `TNODATA` to indicate that the connection has not yet been established. In this way, the function simply initiates the connection establishment procedure by sending a connect request to the destination transport user.

**RETURN VALUES**

**t\_connect()** returns:

- 0 on success.
- 1 on failure and sets **t\_errno** to indicate the error.

**ERRORS**

TACCES	The user does not have permission to use the specified address or options.
TBADADDR	The specified protocol address was in an incorrect format or contained illegal information.
TBADDATA	The amount of user data specified was not within the bounds allowed by the transport provider.
TBADF	The specified file descriptor does not refer to a transport endpoint.
TBADOPT	The specified protocol options were in an incorrect format or contained illegal information.
TBUFOVFLW	The number of bytes allocated for an incoming argument is not sufficient to store the value of that argument. If executed in synchronous mode, the transport provider's state, as seen by the user, changes to <b>T_DATAXFER</b> and the connect indication information to be returned in <i>rcvcall</i> is discarded.
TLOOK	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
TNODATA	<b>T_NDELAY</b> was set, so the function successfully initiated the connection establishment procedure, but did not wait for a response from the remote user.
TNOTSUPPORT	This function is not supported by the underlying transport provider.
TOUTSTATE	The function was issued in the wrong sequence.
TSYSERR	The function failed due to a system error and set <b>errno</b> to indicate the error.

**SEE ALSO**

**intro(3)**, **t\_accept(3N)**, **t\_getinfo(3N)**, **t\_listen(3N)**, **t\_open(3N)**, **t\_optmgmt(3N)**, **t\_rcvconnect(3N)**

*Network Programming*

**NAME**

`t_error` – produce error message

**SYNOPSIS**

```
#include <tiuser.h>

void t_error(errmsg)
char *errmsg;

extern int t_errno;
extern char *t_errlist[ ];
extern int t_nerr;
```

**DESCRIPTION**

`t_error()` produces a message on the standard error output which describes the last error received during a call to a transport function. The argument string *errmsg* is a user-supplied error message that gives context to the error. `t_error()` prints the user-supplied error message followed by a colon and a standard error message for the current error defined in `t_errno`. To simplify variant formatting of messages, the array of message strings `t_errlist` is provided; `t_errno` can be used as an index in this table to get the message string without the NEWLINE. `t_nerr` is the largest message number provided for in the `t_errlist` table.

`t_errno` is only set when an error occurs and is not cleared on successful calls.

**EXAMPLE**

If a `t_connect(3N)` function fails on transport endpoint *fd2* because a bad address was given, the following call might follow the failure:

```
t_error ("t_connect failed on fd2");
```

The diagnostic message to be printed would look like:

```
t_connect failed on fd2: Incorrect transport address format
```

where 'Incorrect transport address format' identifies the specific error that occurred, and 't\_connect failed on fd2' tells the user which function failed on which transport endpoint.

**SEE ALSO**

*Network Programming*

## NAME

`t_free` – free a library structure

## SYNOPSIS

```
#include <tiuser.h>
```

```
int t_free(ptr, struct_type)
```

```
char *ptr;
```

```
int struct_type;
```

## DESCRIPTION

`t_free()` frees memory previously allocated by `t_alloc(3N)`. This function will free memory for the specified structure, and will also free memory for buffers referenced by the structure.

`ptr` points to one of the six structure types described for `t_alloc(3N)`, and `struct_type` identifies the type of that structure which can be one of the following:

<code>T_BIND</code>	<code>struct t_bind</code>
<code>T_CALL</code>	<code>struct t_call</code>
<code>T_OPTMGMT</code>	<code>struct t_optmgmt</code>
<code>T_DIS</code>	<code>struct t_discon</code>
<code>T_UNITDATA</code>	<code>struct t_unitdata</code>
<code>T_UDERROR</code>	<code>struct t_uderr</code>
<code>T_INFO</code>	<code>struct t_info</code>

where each of these structures is used as an argument to one or more transport functions.

`t_free()` checks the `addr`, `opt`, and `udata` fields of the given structure (as appropriate), and frees the buffers pointed to by the `buf` field of the `netbuf` (see `intro(3)`) structure. The `maxlen`, `len`, and `buf` members of the `netbuf` structure are described in `t_accept(3N)`. If `buf` is NULL, `t_free()` will not attempt to free memory. After all buffers are freed, `t_free()` will free the memory associated with the structure pointed to by `ptr`.

Undefined results will occur if `ptr` or any of the `buf` pointers points to a block of memory that was not previously allocated by `t_alloc(3N)`.

## RETURN VALUES

`t_free()` returns:

0        on success.

-1       on failure and sets `t_errno` to indicate the error.

## ERRORS

TSYSERR        The function failed due to a system error and set `errno` to indicate the error.

## SEE ALSO

`intro(3)`, `t_alloc(3N)`

*Network Programming*

## NAME

`t_getinfo` – get protocol-specific service information

## SYNOPSIS

```
#include <tiuser.h>

int t_getinfo(fd, info)
int fd;
struct t_info *info;
```

## DESCRIPTION

`t_getinfo()` returns the current characteristics of the underlying transport protocol associated with file descriptor `fd`. The `info` structure is used to return the same information returned by `t_open(3N)`. `t_getinfo()` enables a transport user to access this information during any phase of communication.

This argument points to a `t_info` structure which contains the following members:

```
long addr;      /* max size of the transport protocol address */
long options;   /* max number of bytes of protocol-specific options */
long tsdu;      /* max size of a transport service data unit (TSDU) */
long etsdu;     /* max size of an expedited transport service data unit (ETSDU) */
long connect;   /* max amount of data allowed on connection establishment
                functions */
long discon;    /* max amount of data allowed on t_snddis and t_rcvdis functions */
long servtype;  /* service type supported by the transport provider */
```

## FIELDS

The values of the fields have the following meanings:

<i>addr</i>	A value greater than or equal to zero indicates the maximum size of a transport protocol address; a value of <code>-1</code> specifies that there is no limit on the address size; and a value of <code>-2</code> specifies that the transport provider does not provide user access to transport protocol addresses.
<i>options</i>	A value greater than or equal to zero indicates the maximum number of bytes of protocol-specific options supported by the provider; a value of <code>-1</code> specifies that there is no limit on the option size; and a value of <code>-2</code> specifies that the transport provider does not support user-settable options.
<i>tsdu</i>	A value greater than zero specifies the maximum size of a transport service data unit (TSDU); a value of zero specifies that the transport provider does not support the concept of TSDU, although it does support the sending of a data stream with no logical boundaries preserved across a connection; a value of <code>-1</code> specifies that there is no limit on the size of a TSDU; and a value of <code>-2</code> specifies that the transfer of normal data is not supported by the transport provider.
<i>etsdu</i>	A value greater than zero specifies the maximum size of an expedited transport service data unit (ETSDU); a value of zero specifies that the transport provider does not support the concept of ETSDU, although it does support the sending of an expedited data stream with no logical boundaries preserved across a connection; a value of <code>-1</code> specifies that there is no limit on the size of an ETSDU; and a value of <code>-2</code> specifies that the transfer of expedited data is not supported by the transport provider.
<i>connect</i>	A value greater than or equal to zero specifies the maximum amount of data that may be associated with connection establishment functions; a value of <code>-1</code> specifies that there is no limit on the amount of data sent during connection establishment; and a value of <code>-2</code> specifies that the transport provider does not allow data to be sent with connection establishment functions.

- discon** A value greater than or equal to zero specifies the maximum amount of data that may be associated with the `t_snddis(3N)` and `t_rcvdis(3N)` functions; a value of `-1` specifies that there is no limit on the amount of data sent with these abortive release functions; and a value of `-2` specifies that the transport provider does not allow data to be sent with the abortive release functions.
- servtype** This field specifies the service type supported by the transport provider, as described below.

If a transport user is concerned with protocol independence, the above sizes may be accessed to determine how large the buffers must be to hold each piece of information. Alternatively, the `t_alloc(3N)` function may be used to allocate these buffers. An error will result if a transport user exceeds the allowed data size on any function. The value of each field may change as a result of option negotiation, and `t_getinfo()` enables a user to retrieve the current characteristics.

#### RETURN VALUES

The *servtype* field of *info* may specify one of the following values on return:

- T\_COTS** The transport provider supports a connection-mode service but does not support the optional orderly release facility.
- T\_COTS\_ORD** The transport provider supports a connection-mode service with the optional orderly release facility.
- T\_CLTS** The transport provider supports a connectionless-mode service. For this service type, `t_open(3N)` will return `-2` for the `etsdu`, `connect`, and `discon` fields.

#### RETURN VALUES

`t_getinfo()` returns 0 on success and `-1` on failure.

#### ERRORS

- TBADF** The specified file descriptor does not refer to a transport endpoint.
- TSYSERR** The function failed due to a system error and set `errno` to indicate the error.

#### SEE ALSO

`t_open(3N)`

*Network Programming*

**NAME**

`t_getstate` – get the current state

**SYNOPSIS**

```
#include <tiuser.h>
```

```
int t_getstate(fd)
```

```
int fd;
```

**DESCRIPTION**

`t_getstate()` returns the current state of the provider associated with the transport endpoint specified by *fd*.

If the provider is undergoing a state transition when `t_getstate()` is called, the function will fail. `t_getstate()` returns the current state on successful completion and `-1` on failure and `t_errno` is set to indicate the error. The current state may be one of the following:

<code>T_UNBND</code>	unbound
<code>T_IDLE</code>	idle
<code>T_OUTCON</code>	outgoing connection pending
<code>T_INCON</code>	incoming connection pending
<code>T_DATAXFER</code>	data transfer
<code>T_OUTREL</code>	outgoing orderly release (waiting for an orderly release indication)
<code>T_INREL</code>	incoming orderly release (waiting for an orderly release request)

**RETURN VALUES**

`t_getstate()` returns:

0 on success.

`-1` on failure and sets `t_errno` to indicate the error.

**ERRORS**

<code>TBADF</code>	The specified file descriptor does not refer to a transport endpoint.
<code>TSTATECHNG</code>	The transport provider is undergoing a state change.
<code>TSYSERR</code>	The function failed due to a system error and set <code>errno</code> to indicate the error.

**SEE ALSO**

`t_open(3N)`

*Network Programming*



**NAME**

`t_listen` – listen for a connect request

**SYNOPSIS**

```
#include <tiuser.h>

int t_listen(fd, call)
int fd;
struct t_call *call;
```

**DESCRIPTION**

`t_listen()` listens for a connect request from a calling transport user. *fd* identifies the local transport endpoint where connect indications arrive, and on return, *call* contains information describing the connect indication. *call* points to a `t_call()` structure which contains the following members:

```
    struct netbuf addr;
    struct netbuf opt;
    struct netbuf udata;
    int sequence;
```

The *maxlen*, *len*, and *buf* members of the *netbuf* structure are described in `t_accept(3N)`. In *call*, *addr* returns the protocol address of the calling transport user, *opt* returns protocol-specific parameters associated with the connect request, *udata* returns any user data sent by the caller on the connect request, and *sequence* is a number that uniquely identifies the returned connect indication. The value of *sequence* enables the user to listen for multiple connect indications before responding to any of them.

Since this function returns values for the *addr*, *opt*, and *udata* fields of *call*, the *maxlen* field of each must be set before issuing the `t_listen()` to indicate the maximum size of the buffer for each.

By default, `t_listen()` executes in synchronous mode and waits for a connect indication to arrive before returning to the user. However, if `T_NDELAY` is set (using `t_open(3N)` or `fcntl()`), `t_listen()` executes asynchronously, reducing to a `poll(2)` for existing connect indications. If none are available, it returns `-1` and sets `t_errno` to `TNODATA`.

**RETURN VALUES**

`t_listen()` returns:

0        on success.

-1       on failure and sets `t_errno` to indicate the error.

**ERRORS**

TBADF	The specified file descriptor does not refer to a transport endpoint.
TBUFOVFLW	The number of bytes allocated for an incoming argument is not sufficient to store the value of that argument. The provider's state, as seen by the user, changes to <code>T_INCON</code> and the connect indication information to be returned in <i>call</i> is discarded.
TLOOK	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
TNODATA	<code>T_NDELAY</code> was set, but no connect indications had been queued.
TNOTSUPPORT	This function is not supported by the underlying transport provider.
TSYSERR	The function failed due to a system error and set <code>errno</code> to indicate the error.

**SEE ALSO****intro(3), t\_accept(3N), t\_bind(3N), t\_connect(3N), t\_open(3N), t\_rcvconnect(3N)***Network Programming*

**NAME**

**t\_look** – look at the current event on a transport endpoint

**SYNOPSIS**

```
#include <tiuser.h>
```

```
int t_look(fd)
```

```
int fd;
```

**DESCRIPTION**

**t\_look()** returns the current event on the transport endpoint specified by *fd*. This function enables a transport provider to notify a transport user of an asynchronous event when the user is issuing functions in synchronous mode. Certain events require immediate notification of the user and are indicated by a specific error, TLOOK, on the current or next function to be executed.

This function also enables a transport user to **poll(2)** a transport endpoint periodically for asynchronous events.

**RETURN VALUES**

Upon success, **t\_look()** returns a value that indicates which of the allowable events has occurred, or returns zero if no event exists. One of the following events is returned:

<b>T_LISTEN</b>	Connection indication received
<b>T_CONNECT</b>	Connect confirmation received
<b>T_DATA</b>	Normal data received
<b>T_EXDATA</b>	Expedited data received
<b>T_DISCONNECT</b>	Disconnect received
<b>T_ERROR</b>	Fatal error indication
<b>T_UDERR</b>	Datagram error indication
<b>T_ORDREL</b>	Orderly release indication

On failure, **-1** is returned and **t\_errno** is set to indicate the error.

**ERRORS**

<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TSYSERR</b>	The function failed due to a system error and set <b>errno</b> to indicate the error.

**SEE ALSO**

**t\_open(3N)**

*Network Programming*

## NAME

`t_open` – establish a transport endpoint

## SYNOPSIS

```
#include <tiuser.h>

int t_open(path, oflag, info)
char *path;
int oflag;
struct t_info *info;
```

## DESCRIPTION

`t_open()` must be called as the first step in the initialization of a transport endpoint. It establishes a transport endpoint by opening a file that identifies a particular transport provider (such as a transport protocol) and returning a file descriptor that identifies that endpoint. For example, opening the file `/dev/tcp` identifies an OSI connection-oriented transport layer protocol as the transport provider. Currently, `/dev/tcp` is the only transport protocol available to `t_open()`.

`path` points to the pathname of the file to open, and `oflag` identifies any open flags (as in `open(2V)`). `t_open()` returns a file descriptor that will be used by all subsequent functions to identify the particular local transport endpoint.

This function also returns various default characteristics of the underlying transport protocol by setting fields in the `t_info` structure pointed to by `info`. `t_info` is defined in `<netli/tiuser.h>` as:

```
struct t_info {
    long addr;      /* size of protocol address */
    long options;  /* size of protocol options */
    long tsdu;     /* size of max transport service data unit */
    long etsdu;    /* size of max expedited tsdu */
    long connect;  /* max data for connection primitives */
    long discon;   /* max data for disconnect primitives */
    long servtype; /* provider service type */
};
```

The fields of this structure have the following values:

- |                |   |
|----------------|---|
| <b>addr</b>    | A value greater than or equal to zero indicates the maximum size of a transport protocol address; a value of <code>-1</code> specifies that there is no limit on the address size; and a value of <code>-2</code> specifies that the transport provider does not provide user access to transport protocol addresses.   |
| <b>options</b> | A value greater than or equal to zero indicates the maximum number of bytes of protocol-specific options supported by the provider; a value of <code>-1</code> specifies that there is no limit on the option size; and a value of <code>-2</code> specifies that the transport provider does not support user-settable options.  |
| <b>tsdu</b>    | A value greater than zero specifies the maximum size of a transport service data unit (TSDU); a value of zero specifies that the transport provider does not support the concept of TSDU, although it does support the sending of a data stream with no logical boundaries preserved across a connection; a value of <code>-1</code> specifies that there is no limit on the size of a TSDU; and a value of <code>-2</code> specifies that the transfer of normal data is not supported by the transport provider.                              |
| <b>etsdu</b>   | A value greater than zero specifies the maximum size of an expedited transport service data unit (ETSDU); a value of zero specifies that the transport provider does not support the concept of ETSDU, although it does support the sending of an expedited data stream with no logical boundaries preserved across a connection; a value of <code>-1</code> specifies that there is no limit on the size of an ETSDU; and a value of <code>-2</code> specifies that the transfer of expedited data is not supported by the transport provider. |

- connect** A value greater than or equal to zero specifies the maximum amount of data that may be associated with connection establishment functions; a value of  $-1$  specifies that there is no limit on the amount of data sent during connection establishment; and a value of  $-2$  specifies that the transport provider does not allow data to be sent with connection establishment functions.
- discon** A value greater than or equal to zero specifies the maximum amount of data that may be associated with the `t_snddis(3N)` and `t_rcvdis(3N)` functions; a value of  $-1$  specifies that there is no limit on the amount of data sent with these abortive release functions; and a value of  $-2$  specifies that the transport provider does not allow data to be sent with the abortive release functions.
- servtype** This field specifies the service type supported by the transport provider. The *servtype* field of *info* may specify one of the following values on return:
- T\_COTS** The transport provider supports a connection-mode service but does not support the optional orderly release facility.
- T\_COTS\_ORD** The transport provider supports a connection-mode service with the optional orderly release facility.
- T\_CLTS** The transport provider supports a connectionless-mode service. For this service type, `t_open()` will return  $-2$  for *etsdu*, *connect*, and *discon*.

A single transport endpoint may support only one of the above services at one time.

If *info* is set to NULL by the transport user, no protocol information is returned by `t_open()`.

If a transport user is concerned with protocol independence, the above sizes may be accessed to determine how large the buffers must be to hold each piece of information. Alternatively, the `t_alloc(3N)` function may be used to allocate these buffers. An error will result if a transport user exceeds the allowed data size on any function.

#### RETURN VALUES

`t_open()` returns a non-negative file descriptor on success. On failure, it returns  $-1$  and sets `t_errno` to indicate the error.

#### ERRORS

**TSYSERR** The function failed due to a system error and set `errno` to indicate the error.

#### SEE ALSO

`open(2V)`, `tcp(4P)`

*Network Programming*

## NAME

`t_optmgmt` – manage options for a transport endpoint

## SYNOPSIS

```
#include <tiuser.h>

int t_optmgmt(fd, req, ret)
int fd;
struct t_optmgmt *req;
struct t_optmgmt *ret;
```

## DESCRIPTION

`t_optmgmt()` enables a transport user to retrieve, verify, or negotiate protocol options with the transport provider. *fd* identifies a bound transport endpoint.

The *req* and *ret* arguments point to a `t_optmgmt()` structure containing the following members:

```
struct netbuf opt;
long flags;
```

The *opt* field identifies protocol options and the *flags* field is used to specify the action to take with those options.

The options are represented by a *netbuff* structure in a manner similar to the address in `t_bind(3N)`. The *maxlen*, *len*, and *buf* members of the *netbuf* structure are described in `t_accept(3N)`. *req* is used to request a specific action of the provider and to send options to the provider. *len* specifies the number of bytes in the options, *buf* points to the options buffer, and *maxlen* has no meaning for the *req* argument. The transport provider may return options and flag values to the user through *ret*. For *ret*, *maxlen* specifies the maximum size of the options buffer and *buf* points to the buffer where the options are to be placed. On return, *len* specifies the number of bytes of options returned. *maxlen* has no meaning for the *req* argument, but must be set in the *ret* argument to specify the maximum number of bytes the options buffer can hold. The actual structure and content of the options is imposed by the transport provider.

The *flags* field of *req* can specify one of the following actions:

<b>T_NEGOTIATE</b>	Enables the user to negotiate the values of the options specified in <i>req</i> with the transport provider. The provider will evaluate the requested options and negotiate the values, returning the negotiated values through <i>ret</i> .
<b>T_CHECK</b>	Enables the user to verify whether the options specified in <i>req</i> are supported by the transport provider. On return, the <i>flags</i> field of <i>ret</i> will have either <b>T_SUCCESS</b> or <b>T_FAILURE</b> set to indicate to the user whether the options are supported. These flags are only meaningful for the <b>T_CHECK</b> request.
<b>T_DEFAULT</b>	Enables a user to retrieve the default options supported by the transport provider into the <i>opt</i> field of <i>ret</i> . In <i>req</i> , the <i>len</i> field of <i>opt</i> must be zero and the <i>buf</i> field may be <b>NULL</b> .

If issued as part of the connectionless-mode service, `t_optmgmt()` may block due to flow control constraints. `t_optmgmt()` will not complete until the transport provider has processed all previously sent data units.

## RETURN VALUES

`t_optmgmt()` returns:

- 0 on success.
- 1 on failure and sets `t_errno` to indicate the error.

**ERRORS**

TACCES	The user does not have permission to negotiate the specified options.
TBADF	The specified file descriptor does not refer to a transport endpoint.
TBADFLAG	An invalid flag was specified.
TBADOPT	The specified protocol options were in an incorrect format or contained illegal information.
TBUFOVFLW	The number of bytes allowed for an incoming argument is not sufficient to store the value of that argument. The information to be returned in <i>ret</i> will be discarded.
TOUTSTATE	The function was issued in the wrong sequence.
TSYSERR	The function failed due to a system error and set <b>errno</b> to indicate the error.

**SEE ALSO**

**intro(3), t\_getinfo(3N), t\_open(3N)**

*Network Programming*

## NAME

`t_rcv` – receive normal or expedited data sent over a connection

## SYNOPSIS

```
int t_rcv(fd, buf, nbytes, flags)
```

```
int fd;
char *buf;
unsigned nbytes;
int *flags;
```

## DESCRIPTION

`t_rcv()` receives either normal or expedited data. *fd* identifies the local transport endpoint through which data will arrive, *buf* points to a receive buffer where user data will be placed, and *nbytes* specifies the size of the receive buffer. *flags* may be set on return from `t_rcv()` and specifies optional flags as described below.

By default, `t_rcv()` operates in synchronous mode and will wait for data to arrive if none is currently available. However, if `T_NDELAY` is set (using `t_open(3N)` or `fcntl()`), `t_rcv()` will execute in asynchronous mode and will fail if no data is available. See `TNODATA` below.

On return from the call, if `T_MORE` is set in *flags* this indicates that there is more data and the current transport service data unit (TSDU) or expedited transport service data unit (ETSDU) must be received in multiple `t_rcv()` calls. Each `t_rcv()` with the `T_MORE` flag set indicates that another `t_rcv()` must follow immediately to get more data for the current TSDU. The end of the TSDU is identified by the return of a `t_rcv()` call with the `T_MORE` flag not set. If the transport provider does not support the concept of a TSDU as indicated in the *info* argument on return from `t_open(3N)` or `t_getinfo(3N)`, the `T_MORE` flag is not meaningful and should be ignored.

On return, the data returned is expedited data if `T_EXPEDITED` is set in *flags*. If the number of bytes of expedited data exceeds *nbytes*, `t_rcv()` will set `T_EXPEDITED` and `T_MORE` on return from the initial call. Subsequent calls to retrieve the remaining ETSDU will not have `T_EXPEDITED` set on return. The end of the ETSDU is identified by the return of a `t_rcv()` call with the `T_MORE` flag not set.

If expedited data arrives after part of a TSDU has been retrieved, receipt of the remainder of the TSDU will be suspended until the ETSDU has been processed. Only after the full ETSDU has been retrieved (`T_MORE` not set) will the remainder of the TSDU be available to the user.

## RETURN VALUES

On success, `t_rcv()` returns the number of bytes received. On failure, it returns `-1`.

## ERRORS

TBADF	The specified file descriptor does not refer to a transport endpoint.
TLOOK	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
TNODATA	<code>T_NDELAY</code> was set, but no data is currently available from the transport provider.
TNOTSUPPORT	This function is not supported by the underlying transport provider.
TSYSERR	The function failed due to a system error and set <code>errno</code> to indicate the error.

## SEE ALSO

`t_open(3N)`, `t_snd(3N)`  
*Network Programming*



## NAME

`t_rcvconnect` – receive the confirmation from a connect request

## SYNOPSIS

```
#include <tiuser.h>

int t_rcvconnect(fd, call)
int fd;
struct t_call *call;
```

## DESCRIPTION

`t_rcvconnect` allows a calling transport user to get the status of a previous connect request. It can be used in conjunction with `t_connect(3N)` to establish a connection in asynchronous mode.

`fd` identifies the local transport endpoint where communication is established. `call` contains information associated with the newly established connection `call` points to a `t_call` structure that contains information associated with the new connection, and is defined in `<netli/tiuser.h>` as:

```
struct t_call {
    struct netbuf addr;
    struct netbuf opt;
    struct netbuf udata;
    int sequence;
};
```

The `maxlen`, `len`, and `buf` members of the `netbuf` structure are described in `t_accept(3N)`. In the `t_call` structure, `addr` returns the protocol address associated with the responding transport endpoint, `opt` presents protocol-specific information associated with the connection, `udata` points to optional user data that may be returned by the destination transport user during connection establishment, and `sequence` has no meaning for this function.

The `maxlen` field of each argument must be set before issuing this function to indicate the maximum buffer size. However, `call` may be NULL, in which case no information is given to the user on return from `t_rcvconnect()`. By default, `t_rcvconnect()` executes synchronously and waits for the connection before returning. On return, the `addr`, `opt`, and `udata` fields reflect values associated with the connection.

If `O_NDELAY` is set (using `t_open(3N)` or `fcntl()`), `t_rcvconnect()` executes asynchronously, reducing to a `poll(2)` request for existing connect confirmations. If none are available, `t_rcvconnect()` fails and returns immediately without waiting for the connection to be established. See `TNODATA` below. `t_rcvconnect()` must be re-issued at a later time to complete the connection establishment phase and retrieve the information returned in `call`.

## RETURN VALUES

`t_rcvconnect()` returns:

- 0        on success.
- 1       on failure and sets `t_errno` to indicate the error.

## ERRORS

- |           |  |
|-----------|--|
| TBADF     | The specified file descriptor does not refer to a transport endpoint.  |
| TBUFOVFLW | The bytes allocated for an incoming argument is sufficient to store the value of that argument and the connect information to be returned in <code>call</code> is discarded. The transport provider's state, as seen by the user, will be changed to <code>DATAXFER</code> . |
| TNODATA   | <code>O_NDELAY</code> was set, but a connect confirmation has not yet arrived.   |
| TLOOK     | An asynchronous event has occurred on this transport connection and requires immediate attention.  |

**TNOTSUPPORT**

This function is not supported by the underlying transport provider.

**TSYSERR**The function failed due to a system error and set `errno` to indicate the error.**SEE ALSO****poll(2), intro(3), t\_accept(3N), t\_bind(3N), t\_connect(3N), t\_listen(3N), t\_open(3N)***Network Programming*

**NAME**

`t_rcvdis` – retrieve information from disconnect

**SYNOPSIS**

```
#include <tiuser.h>

t_rcvdis(fd, discon)
int fd;
struct t_discon *discon;
```

**DESCRIPTION**

`t_rcvdis()` is used to identify the cause of a disconnect, and to retrieve any user data sent with the disconnect. *fd* identifies the local transport endpoint where the connection existed, and *discon* points to a `t_discon` structure defined in `<netli/tiuser.>` as:

```
struct t_discon {
    struct netbuf udata;          /* user data */
    int reason;                  /* reason code */
    int sequence;                /* sequence number */
};
```

The *maxlen*, *len*, and *buf* members of the *netbuf* structure are described in `t_accept(3N)`. *reason* specifies the reason for the disconnect through a protocol-dependent reason code, *udata* identifies any user data that was sent with the disconnect, and *sequence* may identify an outstanding connect indication with which the disconnect is associated. *sequence* is only meaningful when `t_rcvdis()` is issued by a passive transport user who has executed one or more `t_listen(3N)` functions and is processing the resulting connect indications. If a disconnect indication occurs, *sequence* can be used to identify which of the outstanding connect indications is associated with the disconnect.

If a user does not care if there is incoming data and does not need to know the value of *reason* or *sequence*, *discon* may be NULL and any user data associated with the disconnect will be discarded. However, if a user has retrieved more than one outstanding connect indication (using `t_listen(3N)`) and *discon* is NULL, the user will be unable to identify with which connect indication the disconnect is associated.

**RETURN VALUES**

`t_rcvdis()` returns:

- 0 on success.
- 1 on failure and sets `t_errno` to indicate the error.

**ERRORS**

- |             |   |
|-------------|---|
| TBADF       | The specified file descriptor does not refer to a transport endpoint.   |
| TBUFOVFLW   | The number of bytes allocated for incoming data is not sufficient to store the data. The provider's state, as seen by the user, will change to <code>T_IDLE</code> and the disconnect indication information to be returned in <i>discon</i> will be discarded. |
| TNODIS      | No disconnect indication currently exists on the specified transport endpoint.  |
| TNOTSUPPORT | This function is not supported by the underlying transport provider.  |
| TSYSERR     | The function failed due to a system error and set <code>errno</code> to indicate the error.   |

**SEE ALSO**

`intro(3)`, `t_connect(3N)`, `t_listen(3N)`, `t_open(3N)`, `t_snddis(3N)`

*Network Programming*

**NAME**

`t_rcvrel` – acknowledge receipt of an orderly release indication

**SYNOPSIS**

```
#include <tiuser.h>
```

```
int t_rcvrel(fd)
```

```
int fd;
```

**DESCRIPTION**

`t_rcvrel()` acknowledges receipt of an orderly release indication. *fd* identifies the local transport endpoint where the connection exists. After receipt of this indication, the user may not attempt to receive more data because such an attempt will block forever. However, the user may continue to send data over the connection if `t_sndrel(3N)` has not been issued by the user.

`t_rcvrel()` is an optional service of the transport provider, and is only supported if the transport provider returned service type `T_COTS_ORD` on `t_open(3N)` or `t_getinfo(3N)`.

**RETURN VALUES**

`t_rcvrel()` returns:

0        on success.

-1       on failure and sets `t_errno` to indicate the error.

**ERRORS**

<code>TBADF</code>	The specified file descriptor does not refer to a transport endpoint.
<code>TLOOK</code>	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
<code>TNOREL</code>	No orderly release indication currently exists on the specified transport endpoint.
<code>TNOTSUPPORT</code>	This function is not supported by the underlying transport provider.
<code>TSYSERR</code>	The function failed due to a system error and set <code>errno</code> to indicate the error.

**SEE ALSO**

`t_open(3N)`, `t_sndrel(3N)`

*Network Programming*

## NAME

`t_rcvudata` – receive a data unit

## SYNOPSIS

```
#include <tiuser.h>

int t_rcvudata(fd, unitdata, flags)
int fd;
struct t_unitdata *unitdata;
int *flags;
```

## DESCRIPTION

`t_rcvudata()` is used in connectionless mode to receive a data unit from another transport user. *fd* identifies the local transport endpoint through which data will be received, *unitdata* holds information associated with the received data unit, and *flags* is set on return to indicate that the complete data unit was not received. *unitdata* points to a `t_unitdata` structure defined in `<netlli/tiuser.h>` as:

```
struct t_unitdata {
    struct netbuf addr;           /* address      */
    struct netbuf opt;           /* options      */
    struct netbuf udata;         /* user data    */
};
```

The *maxlen*, *len*, and *buf* members of the `netbuf` structure are described in `t_accept(3N)`. The *maxlen* field of *addr*, *opt*, and *udata* must be set before issuing `t_rcvudata()` to indicate the maximum size of the buffer for each.

On return from this call, *addr* specifies the protocol address of the sending user, *opt* identifies protocol-specific options that were associated with this data unit, and *udata* specifies the user data that was received.

By default, `t_rcvudata()` operates in synchronous mode and will wait for a data unit to arrive if none is currently available. However, if `O_NDELAY` is set (using `t_open(3N)` or `fcntl()`), `t_rcvudata()` will execute in asynchronous mode and will fail if no data units are available.

If the buffer defined in the *udata* field of *unitdata* is not large enough to hold the current data unit, the buffer will be filled and `T_MORE` will be set in *flags* on return to indicate that another `t_rcvudata()` should be issued to retrieve the rest of the data unit. Subsequent `t_rcvudata()` call(s) will return zero for the length of the address and options until the full data unit has been received.

## RETURN VALUES

`t_rcvudata()` returns:

- 0        on success.
- 1       on failure and sets `t_errno` to indicate the error.

## ERRORS

- |             |   |
|-------------|---|
| TBADF       | The specified file descriptor does not refer to a transport endpoint.   |
| TBUFOVFLW   | The number of bytes allocated for the incoming protocol address or options is not sufficient to store the information. The unit data information to be returned in <i>unitdata</i> will be discarded. |
| TLOOK       | An asynchronous event has occurred on this transport endpoint and requires immediate attention.   |
| TNODATA     | <code>T_NDELAY</code> was set, but no data units are currently available from the transport provider.   |
| TNOTSUPPORT | This function is not supported by the underlying transport provider.  |
| TSYSERR     | The function failed due to a system error and set <code>errno</code> to indicate the error.   |

## SEE ALSO

`intro(3)`, `t_rcvuderr(3N)`, `t_sndudata(3N)`

## NAME

`t_rcvuderr` – receive a unit data error indication

## SYNOPSIS

```
#include <tiuser.h>

int t_rcvuderr(fd, uderr)
int fd;
struct t_uderr *uderr;
```

## DESCRIPTION

`t_rcvuderr()` is used in connectionless mode to receive information concerning an error on a previously sent data unit, and should only be issued following a unit data error indication. It informs the transport user that a data unit with a specific destination address and protocol options produced an error. *fd* identifies the local transport endpoint through which the error report will be received, and *uderr* points to a `t_uderr()` structure defined in `<netli/tiuser.h>` as:

```
struct t_uderr {
    struct netbuf addr;           /* address */
    struct netbuf opt;          /* options */
    long error;                 /* error code */
};
```

The *maxlen*, *len*, and *buf* members of the *netbuf* structure are described in `t_accept(3N)`. The *maxlen* field of *addr* and *opt* must be set before issuing this function to indicate the maximum size of the buffer for each.

On return from this call, the *addr* structure specifies the destination protocol address of the erroneous data unit, the *opt* structure identifies protocol-specific options that were associated with the data unit, and *error* specifies a protocol-dependent error code.

If the user does not care to identify the data unit that produced an error, *uderr* may be set to NULL and `t_rcvuderr()` will simply clear the error indication without reporting any information to the user.

## RETURN VALUES

`t_rcvuderr()` returns:

- 0        on success.
- 1       on failure and sets `t_errno` to indicate the error.

## ERRORS

- |             |  |
|-------------|--|
| TBADF       | The specified file descriptor does not refer to a transport endpoint.  |
| TBUFOVFLW   | The number of bytes allocated for the incoming protocol address or options is not sufficient to store the information. The unit data error information to be returned in <i>uderr</i> will be discarded. |
| TNOTSUPPORT | This function is not supported by the underlying transport provider.   |
| TNOUDERR    | No unit data error indication currently exists on the specified transport endpoint.  |
| TSYSERR     | The function failed due to a system error and set <code>errno</code> to indicate the error.  |

## SEE ALSO

`intro(3)`, `t_rcvudata(3N)`, `t_sndudata(3N)`

*Network Programming*

**NAME**

`t_snd` – send normal or expedited data over a connection

**SYNOPSIS**

```
#include <tiuser.h>
```

```
int t_snd(fd, buf, nbytes, flags)
```

```
int fd;
```

```
char *buf;
```

```
unsigned nbytes;
```

```
int flags;
```

**DESCRIPTION**

`t_snd()` sends either normal or expedited data. *fd* identifies the local transport endpoint over which data should be sent, *buf* points to the user data, *nbytes* specifies the number of user data bytes to be sent, and *flags* specifies any optional flags described below.

By default, `t_snd()` operates synchronously and may wait if flow control restrictions prevents data acceptance by the local transport provider when the call is made. However, if `O_NDELAY` is set (using `t_open(3N)` or `fcntl()`), `t_snd()` executes asynchronously, and fails immediately if there are flow control restrictions.

On success, `t_snd()` returns the byte total accepted by the transport provider. This normally equals the bytes total specified in *nbytes*. If `O_NDELAY` is set, it is possible that the transport provider will accept only part of the data. In this case, `t_snd()` will set `T_MORE` for the data that was sent (see below) and returns a value less than *nbytes*. If *nbytes* is zero, no data is passed to the provider; `t_snd()` returns zero.

If `T_EXPEDITED` is set in *flags*, the data is sent as expedited data, subject to the interpretations of the transport provider.

`T_MORE` indicates to the transport provider that the transport service data unit (TSDU), or expedited transport service data unit (ETSDU), is being sent through multiple `t_snd()` calls. In these calls, the `T_MORE` flag indicates another `t_snd()` is to follow; the end of TSDU (or ETSDU) is identified by a `t_snd()` call without the `T_MORE` flag. `T_MORE` allows the sender to break up large logical data units, while preserving their boundaries at the other end. The flag does not imply how the data is packaged for transfer below the transport interface. If the transport provider does not support the concept of a TSDU as indicated in the *info* argument on return from `t_open(3N)` or `t_getinfo(3N)`, the `T_MORE` flag is meaningless.

The size of each TSDU or ETSDU must not exceed the transport provider limits as returned by `t_open(3N)` or `t_getinfo(3N)`. Failure to comply results in protocol error `EPROTO`. See `TSYSERR` below.

If `t_snd()` is issued from the `T_IDLE` state, the provider may silently discard the data. If `t_snd()` is issued from any state other than `T_DATAXFER` or `T_IDLE` the provider generates a `EPROTO` error.

**RETURN VALUES**

On success, `t_snd()` returns the number of bytes accepted by the transport provider. On failure, it returns `-1` and sets `t_errno` to indicate the error.

**ERRORS**

<code>TBADF</code>	The specified file descriptor does not refer to a transport endpoint.
<code>TFLOW</code>	<code>O_NDELAY</code> was set, but the flow control mechanism prevented the transport provider from accepting data at this time.
<code>TNOTSUPPORT</code>	This function is not supported by the underlying transport provider.
<code>TSYSERR</code>	The function failed due to a system error and set <code>errno</code> to indicate the error.

**SEE ALSO**

**t\_open(3N), t\_rcv(3N)**

*Network Programming*



## NAME

`t_snddis` – send user-initiated disconnect request

## SYNOPSIS

```
#include <tiuser.h>

int t_snddis(fd, call)
int fd;
struct t_call *call;
```

## DESCRIPTION

`t_snddis()` is used to initiate an abortive release on an already established connection or to reject a connect request. *fd* identifies the local transport endpoint of the connection, and *call* specifies information associated with the abortive release. *call* points to a `t_call()` structure which is defined in `<netllie/tiuser.h>` as:

```
struct t_call {
    struct netbuf addr;           /* address          */
    struct netbuf opt;           /* options          */
    struct netbuf udata;        /* user data        */
    int sequence;               /* sequence number  */
};
```

The *maxlen*, *len*, and *buf* members of the *netbuf* structure are described in `t_accept(3N)`. The values in *call* have different semantics, depending on the context of the call to `t_snddis()`. When rejecting a connect request, *call* must be non-NULL and contain a valid value of *sequence* to uniquely identify the rejected connect indication to the transport provider. The *addr* and *opt* fields of *call* are ignored. In all other cases, *call* need only be used when data is being sent with the disconnect request. The *addr*, *opt*, and *sequence* fields of the `t_call()` structure are ignored. If the user does not wish to send data to the remote user, the value of *call* may be NULL. *udata* specifies the user data to be sent to the remote user. The amount of user data must not exceed the limits supported by the transport provider as returned by `t_open(3N)` or `t_getinfo(3N)`. If the *len* field of *udata* is zero, no data will be sent to the remote user.

## RETURN VALUES

`t_snddis()` returns:

- 0        on success.
- 1       on failure and sets `t_errno` to indicate the error.

## ERRORS

- |             |  |
|-------------|--|
| TBADDATA    | The amount of user data specified was not within the bounds allowed by the transport provider. The transport provider's outgoing queue will be flushed, so data may be lost. |
| TBADF       | The specified file descriptor does not refer to a transport endpoint.  |
| TBADSEQ     | An invalid sequence number was specified. The transport provider's outgoing queue will be flushed, so data may be lost.  |
| TLOOK       | A NULL call structure was specified when rejecting a connect request. The transport provider's outgoing queue will be flushed, so data may be lost.                          |
| TLOOK       | An asynchronous event has occurred on this transport endpoint and requires immediate attention.  |
| TNOTSUPPORT | This function is not supported by the underlying transport provider.   |
| TOUTSTATE   | The function was issued in the wrong sequence. The transport provider's outgoing queue may be flushed, so data may be lost.  |
| TSYSERR     | The function failed due to a system error and set <code>errno</code> to indicate the error.  |

**SEE ALSO**

**intro(3), t\_connect(3N), t\_getinfo(3N), t\_listen(3N), t\_open(3N)**

*Network Programming*

**NAME**

`t_sndrel` – initiate an orderly release

**SYNOPSIS**

```
#include <tiuser.h>
```

```
int t_sndrel(fd)
```

```
int fd;
```

**DESCRIPTION**

`t_sndrel()` initiates an orderly release of a transport connection and indicates to the transport provider that the transport user has no more data to send. *fd* identifies the local transport endpoint where the connection exists. After issuing `t_sndrel()`, the user may not send any more data over the connection. However, a user may continue to receive data if an orderly release indication has been received.

`t_sndrel()` is an optional service of the transport provider, and is only supported if the transport provider returned service type `T_COTS_ORD` on `t_open(3N)` or `t_getinfo(3N)`.

**RETURN VALUES**

`t_sndrel()` returns:

0 on success.

-1 on failure and sets `t_errno` to indicate the error.

**ERRORS**

`TBADF` The specified file descriptor does not refer to a transport endpoint.

`TFLOW` `O_NDELAY` was set, but the flow control mechanism prevented the transport provider from accepting the function at this time.

`TNOTSUPPORT` This function is not supported by the underlying transport provider.

`TSYSERR` The function failed due to a system error and set `errno` to indicate the error.

**SEE ALSO**

`t_open(3N)`, `t_rcvrel(3N)`

*Network Programming*

## NAME

t\_sndudata – send a data unit

## SYNOPSIS

```
#include <tiuser.h>

int t_sndudata(fd, unitdata)
int fd;
struct t_unitdata *unitdata;
```

## DESCRIPTION

t\_sndudata() is used in connectionless mode to send a data unit to another transport user. *fd* identifies the local transport endpoint through which data will be sent, and *unitdata* points to a t\_unitdata structure defined in <netlli/tiuser.h> as:

```
struct t_unitdata {
    struct netbuf addr;           /* address      */
    struct netbuf opt;          /* options     */
    struct netbuf udata;        /* user data   */
};
```

The *maxlen*, *len*, and *buf* members of the *netbuf* structure are described in t\_accept(3N). In *unitdata*, *addr* specifies the protocol address of the destination user, *opt* identifies protocol-specific options that the user wants associated with this request, and *udata* specifies the user data to be sent. The user may choose not to specify what protocol options are associated with the transfer by setting the *len* field of *opt* to 0. In this case, the provider may use default options.

If the *len* field of *udata* is 0, no data unit will be passed to the transport provider; t\_sndudata() will not send zero-length data units.

By default, t\_sndudata() operates in synchronous mode and may wait if flow control restrictions prevent the data from being accepted by the local transport provider at the time the call is made. However, if T\_NDELAY is set (using t\_open(3N) or fcntl()), t\_sndudata() will execute in asynchronous mode and will fail under such conditions.

If t\_sndudata() is issued from an invalid state, or if the amount of data specified in *udata* exceeds the TSDU size as returned by t\_open() or t\_getinfo(3N), the provider will generate an EPROTO protocol error. See TSYSErr below.

## RETURN VALUES

t\_sndudata() returns:

- 0 on success.
- 1 on failure and sets t\_errno to indicate the error.

## ERRORS

- |             |   |
|-------------|---|
| TBADF       | The specified file descriptor does not refer to a transport endpoint.   |
| TFLOW       | T_NDELAY was set, but the flow control mechanism prevented the transport provider from accepting data at this time. |
| TNOTSUPPORT | This function is not supported by the underlying transport provider.  |
| TSYSERR     | The function failed due to a system error and set errno to indicate the error.                                      |

## SEE ALSO

intro(3), t\_rcvudata(3N), t\_rcvuderr(3N)

*Network Programming*

**NAME**

`t_sync` – synchronize transport library

**SYNOPSIS**

```
#include <tiuser.h>

int t_sync(fd)
int fd;
```

**DESCRIPTION**

For the transport endpoint specified by *fd*, `t_sync()` synchronizes the data structures managed by the transport library with information from the underlying transport provider. In doing so, it can convert a raw file descriptor (obtained using `open(2V)`, `dup(2V)`, or as a result of a `fork(2V)` and `execve(2V)`) to an initialized transport endpoint, assuming that file descriptor referenced a transport provider. `t_sync()` also allows two cooperating processes to synchronize their interaction with a transport provider.

For example, if a process *forks* a new process and issues an *exec*, the new process must issue a `t_sync()` to build the private library data structure associated with a transport endpoint and to synchronize the data structure with the relevant provider information.

It is important to remember that the transport provider treats all users of a transport endpoint as a single user. If multiple processes are using the same endpoint, they should coordinate their activities so as not to violate the state of the provider. `t_sync()` returns the current state of the provider to the user, thereby enabling the user to verify the state before taking further action. This coordination is only valid among cooperating processes; it is possible that a process or an incoming event could change the provider's state *after* a `t_sync()` is issued.

If the provider is undergoing a state transition when `t_sync()` is called, the function will fail.

**RETURN VALUES**

`t_sync()` returns `-1` on failure. Upon success, the state of the transport provider is returned; it may be one of the following:

<code>T_IDLE</code>	idle
<code>T_OUTCON</code>	outgoing connection pending
<code>T_INCON</code>	incoming connection pending
<code>T_DATAXFER</code>	data transfer
<code>T_OUTREL</code>	outgoing orderly release (waiting for an orderly release indication)
<code>T_INREL</code>	incoming orderly release (waiting for an orderly release request)
<code>T_UNBND</code>	unbound

**ERRORS**

<code>TBADF</code>	The specified file descriptor is a valid open file descriptor but does not refer to a transport endpoint.
<code>TSTATECHNG</code>	The transport provider is undergoing a state change.
<code>TSYSERR</code>	The function failed due to a system error and set <code>errno</code> to indicate the error.

**SEE ALSO**

`dup(2V)`, `execve(2V)`, `fork(2V)`, `open(2V)`

*Network Programming*

**NAME**

`t_unbind` – disable a transport endpoint

**SYNOPSIS**

```
#include <tiuser.h>
```

```
int t_unbind(fd)
```

```
int fd;
```

**DESCRIPTION**

`t_unbind()` disables the transport endpoint specified by `fd` which was previously bound by `t_bind(3N)`. On completion of this call, no further data or events destined for this transport endpoint will be accepted by the transport provider.

**RETURN VALUES**

`t_unbind()` returns:

0        on success.

-1       on failure and sets `t_errno` to indicate the error.

**ERRORS**

TBADF	The specified file descriptor does not refer to a transport endpoint.
TLOOK	An asynchronous event has occurred on this transport endpoint.
TOUTSTATE	The function was issued in the wrong sequence.
TSYSERR	The function failed due to a system error and set <code>errno</code> to indicate the error.

**SEE ALSO**

`t_bind(3N)`

*Network Programming*

**NAME**

tcgetpgrp, tcsetpgrp – get, set foreground process group ID

**SYNOPSIS**

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

```
pid_t tcgetpgrp(fd)
```

```
int fd;
```

```
int tcsetpgrp(fd, pgrp_id)
```

```
int fd;
```

```
pid_t pgrp_id;
```

**DESCRIPTION**

**tcgetpgrp()** returns the value of the process group ID of the foreground process group associated with the terminal (see NOTES). **tcgetpgrp()** is allowed from a process that is a member of a background process group; however, the information may be subsequently changed by a process that is a member of a foreground process group.

If the process has a controlling terminal, **tcsetpgrp()** sets the foreground process group ID associated with the terminal to *pgrp\_id*. The file associated with *fd* must be the controlling terminal and must be currently associated with the session of the calling process. The value of *pgrp\_id* must match a process group ID of a process in the same session as the calling process.

**RETURN VALUES**

On success, **tcgetpgrp()** returns the process group ID of the foreground process group associated with the terminal. On failure, it returns  $-1$  and sets **errno** to indicate the error.

**tcsetpgrp()** returns:

0 on success.

$-1$  on failure and sets **errno** to indicate the error.

**ERRORS**

If any of the following conditions occur, **tcgetpgrp()** sets **errno** to:

EBADF *fd* is not a valid file descriptor.

ENOSYS **tcgetpgrp()** is not supported in this implementation.

ENOTTY The calling process does not have a controlling terminal.

The file is not the controlling terminal.

If any of the following conditions occur, **tcsetpgrp()** sets **errno** to:

EBADF *fd* is not a valid file descriptor.

EINVAL The value of *pgrp\_id* is not a valid process group ID.

ENOTTY The calling process does not have a controlling terminal.

The file is not the controlling terminal.

The controlling terminal is no longer associated with the session of the calling process.

EPERM The value of *pgrp\_id* is a valid process group ID, but does not match the process group ID of a process in the same session as the calling process.

**SEE ALSO**

setpgid(2V), setsid(2V)

## NOTES

For `tcgetpgrp()` and `tcsetpgrp()` to behave as described above, `{_POSIX_JOB_CONTROL}` must be in effect (see `sysconf(2V)`). `{_POSIX_JOB_CONTROL}` is always in effect on SunOS systems, but for portability, applications should call `sysconf()` to determine whether `{_POSIX_JOB_CONTROL}` is in effect for the current system.

If `{_POSIX_JOB_CONTROL}` is not defined on a system conforming to *IEEE Std 1003.1-1988* either `tcgetpgrp()` and `tcsetpgrp()` behave as described above, or `tcgetpgrp()` and `tcsetpgrp()` fail.



## NAME

termcap, tgetent, tgetnum, tgetflag, tgetstr, tgoto, tputs – terminal independent operation routines

## SYNOPSIS

```

char PC;
char *BC;
char *UP;
short ospeed;

tgetent(bp, name)
char *bp, *name;

tgetnum(id)
char *id;

tgetflag(id)
char *id;

char *
tgetstr(id, area)
char *id, **area;

char *
tgoto(cm, destcol, destline)
char *cm;

tputs(cp, affent, outc)
register char *cp;
int affent;
int (*outc)();

```

## DESCRIPTION

These functions extract and use capabilities from the terminal capability data base **termcap(5)**. These are low level routines; see **curses(3V)** for a higher level package.

**tgetent()** extracts the entry for terminal *name* into the *bp* buffer, with the current size of the tty (usually a window). This allows pre-SunWindows programs to run in a window of arbitrary size. *bp* should be a character buffer of size 1024 and must be retained through all subsequent calls to **tgetnum()**, **tgetflag()**, and **tgetstr()**. **tgetent()** returns -1 if it cannot open the **termcap()** file, 0 if the terminal name given does not have an entry, and 1 if all goes well. It will look in the environment for a **TERMCAP** variable. If found, and the value does not begin with a slash, and the terminal type *name* is the same as the environment string **TERM**, the **TERMCAP** string is used instead of reading the **termcap** file. If it does begin with a slash, the string is used as a path name rather than **/etc/termcap**. This can speed up entry into programs that call **tgetent**, as well as to help debug new terminal descriptions or to make one for your terminal if you cannot write the file **/etc/termcap**. Note: if the window size changes, the "lines" and "columns" entries in *bp* are no longer correct. See the *SunView Programmer's Guide* for details regarding [how to handle] this.

**tgetnum()** gets the numeric value of capability **ID**, returning -1 if is not given for the terminal. **tgetflag()** returns 1 if the specified capability is present in the terminal's entry, 0 if it is not. **tgetstr()** gets the string value of capability **ID**, placing it in the buffer at *area*, advancing the *area* pointer. It decodes the abbreviations for this field described in **termcap(5)**, except for cursor addressing and padding information. **tgetstr()** returns the string pointer if successful. Otherwise it returns zero.

**tgoto()** returns a cursor addressing string decoded from *cm* to go to column *destcol* in line *destline*. It uses the external variables **UP** (from the **up** capability) and **BC** (if **bc** is given rather than **bs**) if necessary to avoid placing **\n**, **^D** or **^@** in the returned string. (Programs which call **tgoto()** should be sure to turn off the **XTABS** bit(s), since **tgoto()** may now output a tab. Note: programs using **termcap()** should in general turn off **XTABS** anyway since some terminals use **^I** (CTRL-I) for other functions, such as nondestructive space.) If a **%** sequence is given which is not understood, then **tgoto()** returns **OOPS**.

**tputs()** decodes the leading padding information of the string *cp*; *affcnt* gives the number of lines affected by the operation, or 1 if this is not applicable, *outc* is a routine which is called with each character in turn. The external variable *ospeed* should contain the encoded output speed of the terminal as described in **tty(4)**. The external variable **PC** should contain a pad character to be used (from the **pc** capability) if a **NULL** (**^@**) is inappropriate.

**FILES**

**/usr/lib/libtermcap.a** -termcap library  
**/etc/termcap** data base

**SEE ALSO**

**ex(1)**, **curses(3V)**, **tty(4)**, **termcap(5)**

## NAME

termios, tcgetattr, tcsetattr, tcsendbreak, tcdrain, tcflush, tcflow, cfgetospeed, cfgetispeed, cfsetispeed, cfsetospeed – get and set terminal attributes, line control, get and set baud rate, get and set terminal foreground process group ID

## SYNOPSIS

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

int tcgetattr(fd, termios_p)
int fd;
struct termios *termios_p;

int tcsetattr(fd, optional_actions, termios_p)
int fd;
int optional_actions;
struct termios *termios_p;

int tcsetattr(fd, duration)
int fd;
int duration;

int tcdrain(fd)
int fd;

int tcflush(fd, queue_selector)
int fd;
int queue_selector;

int tcflow(fd, action)
int fd;
int action;

speed_t cfgetospeed(termios_p)
struct termios *termios_p;

int cfsetospeed(termios_p, speed)
struct termios *termios_p;
speed_t speed;

speed_t cfgetispeed(termios_p)
struct termios *termios_p;

int cfsetispeed(termios_p, speed)
struct termios *termios_p;
speed_t speed;

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

## DESCRIPTION

The termios functions describe a general terminal interface that is provided to control asynchronous communications ports. A more detailed overview of the terminal interface can be found in **termio(4)**. That section also describes an **ioctl()** interface that can be used to access the same functionality. However, the function interface described here is the preferred user interface.

Many of the functions described here have a *termios\_p* argument that is a pointer to a **termios** structure. This structure contains the following members:

```

tcflag_t   c_iflag;           /* input modes */
tcflag_t   c_oflag;          /* output modes */
tcflag_t   c_cflag;         /* control modes */
tcflag_t   c_lflag;         /* local modes */
cc_t       c_cc[NCCS];       /* control chars */

```

These structure members are described in detail in `termio(4)`.

`tcgetattr()` gets the parameters associated with the object referred by *fd* and stores them in the `termios` structure referenced by *termios\_p*. This function may be invoked from a background process; however, the terminal attributes may be subsequently changed by a foreground process.

`tcsetattr()` sets the parameters associated with the terminal (unless support is required from the underlying hardware that is not available) from the `termios` structure referred to by *termios\_p* as follows:

- If *optional\_actions* is `TCSANOW`, the change occurs immediately.
- If *optional\_actions* is `TCSADRAIN`, the change occurs after all output written to *fd* has been transmitted. This function should be used when changing parameters that affect output.
- If *optional\_actions* is `TCSAFLUSH`, the change occurs after all output written to the object referred by *fd* has been transmitted, and all input that has been received but not read will be discarded before the change is made.

The symbolic constants for the values of *optional\_actions* are defined in `<sys/termios.h>`.

If the terminal is using asynchronous serial data transmission, `tcsendbreak()` transmits a continuous stream of zero-valued bits for a specific duration. If *duration* is zero, it transmits zero-valued bits for at least 0.25 seconds, and not more than 0.5 seconds. If *duration* is not zero, it sends zero-valued bits for *duration*\**N* seconds, where *N* is at least 0.25, and not more than 0.5.

If the terminal is not using asynchronous serial data transmission, `tcsendbreak()` returns without taking any action.

`tcdrain()` waits until all output written to the object referred to by *fd* has been transmitted.

`tcflush()` discards data written to the object referred to by *fd* but not transmitted, or data received but not read, depending on the value of *queue\_selector*:

- If *queue\_selector* is `TCIFLUSH`, it flushes data received but not read.
- If *queue\_selector* is `TCOFLUSH`, it flushes data written but not transmitted.
- If *queue\_selector* is `TCIOFLUSH`, it flushes both data received but not read, and data written but not transmitted.

The symbolic constants for the values of *queue\_selector* and *action* are defined in `termios.h`.

The default on open of a terminal file is that neither its input nor its output is suspended.

`tcflow()` suspends transmission or reception of data on the object referred to by *fd*, depending on the value of *actions*:

- If *action* is `TCOOFF`, it suspends output.
- If *action* is `TCOON`, it restarts suspended output.
- If *action* is `TCIOFF`, the system transmits a STOP character, which stops the terminal device from transmitting data to the system. (See `termio(4)`.)
- If *action* is `TCION`, the system transmits a START character, which starts the terminal device transmitting data to the system. (See `termio(4)`.)

The baud rate functions are provided for getting and setting the values of the input and output baud rates in the `termios` structure. The effects on the terminal device described below do not become effective until `tcsetattr()` is successfully called.

The input and output baud rates are stored in the `termios` structure. The values shown in the table are supported. The names in this table are defined in `termios.h`

Name	Description	Name	Description
<b>B0</b>	Hang up	<b>B600</b>	600 baud
<b>B50</b>	50 baud	<b>B1200</b>	1200 baud
<b>B75</b>	75 baud	<b>B1800</b>	1800 baud
<b>B110</b>	110 baud	<b>B2400</b>	2400 baud
<b>B134</b>	134.5 baud	<b>B4800</b>	4800 baud
<b>B150</b>	150 baud	<b>B9600</b>	9600 baud
<b>B200</b>	200 baud	<b>B19200</b>	19200 baud
<b>B300</b>	300 baud	<b>B38400</b>	38400 baud

`cfgetospeed()` returns the output baud rate stored in the `termios` structure pointed to by `termios_p`.

`cfsetospeed()` sets the output baud rate stored in the `termios` structure pointed to by `termios_p` to `speed`. The zero baud rate, **B0**, is used to terminate the connection. If **B0** is specified, the modem control lines shall no longer be asserted. Normally, this will disconnect the line.

If the input baud rate is set to zero, the input baud rate will be specified by the value of the output baud rate.

`cfgetispeed()` returns the input baud rate stored in the `termios` structure.

`cfsetispeed()` sets the input baud rate stored in the `termios` structure to `speed`.

#### RETURN VALUES

`cfgetispeed()` returns the input baud rate stored in the `termios` structure.

`cfgetospeed()` returns the output baud rate stored in the `termios` structure.

`cfsetispeed()` and `cfsetospeed()` return:

0 on success.

-1 on failure and sets `errno` to indicate the error.

All other functions return:

0 on success.

-1 on failure and set `errno` to indicate the error.

#### ERRORS

**EBADF** The `fd` argument is not a valid file descriptor.

**ENOTTY** The file associated with `fd` is not a terminal.

`tcsetattr()` may set `errno` to:

**EINVAL** The `optional_actions` argument is not a proper value.

An attempt was made to change an attribute represented in the `termios` structure to an unsupported value.

`tcsendbreak()` may set `errno` to:

**EINVAL** The device does not support `tcsendbreak()`.

`tcdrain()` may set `errno` to:

**EINTR** A signal interrupted `tcdrain()`.

**EINVAL** The device does not support `tcdrain()`.

`tcflush()` may set `errno` to:

**EINVAL** The device does not support `tcflush()`.

The `queue_selector` argument is not a proper value.

**tcflow()** may set **errno** to:

**EINVAL**           The device does not support **tcflow()**.  
                  The *action* argument is not a proper value.

**tcsetattr()** may set **errno** to:

**EAGAIN**           There is insufficient memory available to copy in the arguments.  
**EBADF**           *fd* is not a valid descriptor.  
**EFAULT**          Some part of the structure pointed to by *termios\_p* is outside the process's allocated address space.  
**EINVAL**          *optional\_actions* is not valid.  
**EIO**             The calling process is a background process.  
**ENOTTY**          *fd* does not refer to a terminal device.  
**ENXIO**           The terminal referred to by *fd* is hung up.

**cfsetispeed()** and **cfsetospeed()** may set **errno** to:

**EINVAL**          *speed* is greater than B38400 or less than 0.

**SEE ALSO**

**setpgid(2V)**, **setsid(2V)**, **termio(4)**

## NAME

time, ftime – get date and time

## SYNOPSIS

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

time_t time(tloc)
time_t *tloc;

#include <sys/timeb.h>

int ftime(tp)
struct timeb *tp;
```

## DESCRIPTION

**time()** returns the time since 00:00:00 GMT, Jan. 1, 1970, measured in seconds.

If *tloc* is non-NULL, the return value is also stored in the location to which *tloc* points.

**ftime()** fills in a structure pointed to by *tp*, as defined in **<sys/timeb.h>**:

```
struct timeb
{
    time_t    time;
    unsigned short millitm;
    short    timezone;
    short    dstflag;
};
```

The structure contains the time since the epoch in seconds, up to 1000 milliseconds of more-precise interval, the local time zone (measured in minutes of time westward from Greenwich), and a flag that, if nonzero, indicates that Daylight Saving time applies locally during the appropriate part of the year.

## RETURN VALUES

**time()** returns the value of time on success. On failure, it returns **(time\_t) -1**.

On success, **ftime()** returns no useful value. On failure, it returns **-1**.

## SEE ALSO

**date(1V)**, **gettimeofday(2)**, **ctime(3V)**

**NAME**

times – get process times

**SYNOPSIS**

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

int times(buffer)
struct tms *buffer;
```

**SYSTEM V SYNOPSIS**

```
clock_t times(buffer)
struct tms *buffer;
```

**DESCRIPTION**

This interface is obsoleted by `getrusage(2)`.

`times()` returns time-accounting information for the current process and for the terminated child processes of the current process. All times are in 1/HZ seconds, where HZ is 60.

*buffer* points to the following structure:

```
struct tms {
    clock_t tms_utime;           /* user time */
    clock_t tms_stime;           /* system time */
    clock_t tms_cutime;          /* user time, children */
    clock_t tms_cstime;          /* system time, children */
};
```

This information comes from the calling process and each of its terminated child processes for which it has executed a `wait(2V)`.

`tms_utime` is the CPU time used while executing instructions in the user space of the calling process.

`tms_stime` is the CPU time used by the system on behalf of the calling process.

`tms_cutime` is the sum of the `tms_utimes` and `tms_cutimes` of the child processes.

`tms_cstime` is the sum of the `tms_stimes` and `tms_cstimes` of the child processes.

**RETURN VALUES**

`times()` returns:

```
0      on success.
-1     on failure.
```

**SYSTEM V RETURN VALUES**

Upon successful completion, `times()` returns the elapsed real time, in 60ths of a second, since an arbitrary point in the past. This point does not change from one invocation of `times()` to another within the same process. On failure, `times()` returns `(clock_t) -1`.

**SEE ALSO**

`time(1V)`, `getrusage(2)`, `wait(2V)`, `time(3V)`



**NAME**

timezone – get time zone name given offset from GMT

**SYNOPSIS**

**char \*timezone(zone, dst)**

**DESCRIPTION**

**timezone()** attempts to return the name of the time zone associated with its first argument, which is measured in minutes westward from Greenwich. If the second argument is 0, the standard name is used, otherwise the Daylight Savings Time version. If the required name does not appear in a table built into the routine, the difference from GMT is produced; for instance, in Afghanistan '**timezone(-(60\*4+30), 0)**' is appropriate because it is 4:30 ahead of GMT and the string **GMT+4:30** is produced.

Note: the offset westward from Greenwich and an indication of whether Daylight Savings Time is in effect may not be sufficient to determine the name of the time zone, as the name may differ between different locations in the same time zone. Instead of using **timezone()** to determine the name of the time zone for a given time, that time should be converted to a '**struct tm**' using **localtime()** (see **ctime(3V)**) and the *tm\_zone* field of that structure should be used. **timezone()** is retained for compatibility with existing programs.

**SEE ALSO**

**ctime(3V)**

**NAME**

**tmpfile** – create a temporary file

**SYNOPSIS**

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

```
FILE *tmpfile()
```

**DESCRIPTION**

**tmpfile()** creates a temporary file using a name generated by **tmpnam(3S)**, and returns a corresponding **FILE** pointer. If the file cannot be opened, an error message is printed using **perror(3)**, and a **NULL** pointer is returned. The file will automatically be deleted when the process using it terminates. The file is opened for update ("w+").

**SEE ALSO**

**creat(2V)**, **unlink(2V)**, **fopen(3V)**, **mktemp(3)**, **perror(3)**, **tmpnam(3S)**

**NAME**

`tmpnam`, `tmpnam` – create a name for a temporary file

**SYNOPSIS**

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

```
char *tmpnam (s)
```

```
char *s;
```

```
char *tmpnam (dir, pfx)
```

```
char *dir, *pfx;
```

**DESCRIPTION**

These functions generate file names that can safely be used for a temporary file.

`tmpnam()` always generates a file name using the path-prefix defined as `P_tmpdir` in the `<stdio.h>` header file. If `s` is `NULL`, `tmpnam()` leaves its result in an internal static area and returns a pointer to that area. The next call to `tmpnam()` will destroy the contents of the area. If `s` is not `NULL`, it is assumed to be the address of an array of at least `L_tmpnam` bytes, where `L_tmpnam` is a constant defined in `<stdio.h>`; `tmpnam()` places its result in that array and returns `s`.

`tmpnam()` allows the user to control the choice of a directory. The argument `dir` points to the name of the directory in which the file is to be created. If `dir` is `NULL` or points to a string which is not a name for an appropriate directory, the path-prefix defined as `P_tmpdir` in the `<stdio.h>` header file is used. If that directory is not accessible, `/tmp` will be used as a last resort. This entire sequence can be up-staged by providing an environment variable `TMPDIR` in the user's environment, whose value is the name of the desired temporary-file directory.

Many applications prefer their temporary files to have certain favorite initial letter sequences in their names. Use the `pfx` argument for this. This argument may be `NULL` or point to a string of up to five characters to be used as the first few characters of the temporary-file name.

`tmpnam()` uses `malloc()` to get space for the constructed file name, and returns a pointer to this area. Thus, any pointer value returned from `tmpnam()` may serve as an argument to `free` (see `malloc(3V)`). If `tmpnam()` cannot return the expected result for any reason, that is, `malloc()` failed, or none of the above mentioned attempts to find an appropriate directory was successful, a `NULL` pointer will be returned.

**NOTES**

These functions generate a different file name each time they are called.

Files created using these functions and either `fopen()` or `creat()` are temporary only in the sense that they reside in a directory intended for temporary use, and their names are unique. It is the user's responsibility to use `unlink(2V)` to remove the file when its use is ended.

**SEE ALSO**

`creat(2V)`, `unlink(2V)`, `fopen(3V)`, `malloc(3V)`, `mktemp(3)`, `tmpfile(3S)`

**BUGS**

If called more than 17,576 times in a single process, these functions will start recycling previously used names.

Between the time a file name is created and the file is opened, it is possible for some other process to create a file with the same name. This can never happen if that other process is using these functions or `mktemp()`, and the file names are chosen so as to render duplication by other means unlikely.

## NAME

`tsearch`, `tfind`, `tdelete`, `twalk` – manage binary search trees

## SYNOPSIS

```
#include <search.h>

char *tsearch((char *) key, (char **) rootp, compar)
int (*compar)( );

char *tfind((char *) key, (char **) rootp, compar)
int (*compar)( );

char *tdelete((char *) key, (char **) rootp, compar)
int (*compar)( );

void twalk((char *) root, action)
void (*action)( );
```

## DESCRIPTION

`tsearch()`, `tfind()`, `tdelete()`, and `twalk()` are routines for manipulating binary search trees. They are generalized from Knuth (6.2.2) Algorithms T and D. All comparisons are done with a user-supplied routine. This routine is called with two arguments, the pointers to the elements being compared. It returns an integer less than, equal to, or greater than 0, according to whether the first argument is to be considered less than, equal to or greater than the second argument. The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

`tsearch()` is used to build and access the tree. *key* is a pointer to a datum to be accessed or stored. If there is a datum in the tree equal to *\*key* (the value pointed to by *key*), a pointer to this found datum is returned. Otherwise, *\*key* is inserted, and a pointer to it returned. Only pointers are copied, so the calling routine must store the data. *rootp* points to a variable that points to the root of the tree. A NULL value for the variable pointed to by *rootp* denotes an empty tree; in this case, the variable will be set to point to the datum which will be at the root of the new tree.

Like `tsearch()`, `tfind()` will search for a datum in the tree, returning a pointer to it if found. However, if it is not found, `tfind()` will return a NULL pointer. The arguments for `tfind()` are the same as for `tsearch()`.

`tdelete()` deletes a node from a binary search tree. The arguments are the same as for `tsearch()`. The variable pointed to by *rootp* will be changed if the deleted node was the root of the tree. `tdelete()` returns a pointer to the parent of the deleted node, or a NULL pointer if the node is not found.

`twalk()` traverses a binary search tree. *root* is the root of the tree to be traversed. (Any node in a tree may be used as the root for a walk below that node.) *action* is the name of a routine to be invoked at each node. This routine is, in turn, called with three arguments. The first argument is the address of the node being visited. The second argument is a value from an enumeration data type `typedef enum { preorder, postorder, endorder, leaf } VISIT;` (defined in the `<search.h>` header file), depending on whether this is the first, second or third time that the node has been visited (during a depth-first, left-to-right traversal of the tree), or whether the node is a leaf. The third argument is the level of the node in the tree, with the root being level zero.

The pointers to the key and the root of the tree should be of type pointer-to-element, and cast to type pointer-to-pointer-to-character. Similarly, although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

## EXAMPLES

The following code reads in strings and stores structures containing a pointer to each string and a count of its length. It then walks the tree, printing out the stored strings and their lengths in alphabetical order.

```

#include <search.h>
#include <stdio.h>

void twalk();
char *tsearch();

struct node {          /* pointers to these are stored in the tree */
    char *string;
    int count;
};

#define MAXNODES      12
#define MAXSTRING     100
#define MINSTRING     3          /* char, newline, eos */

char string_space[MAXSTRING];    /* space to store strings */
struct node node_space[MAXNODES]; /* nodes to store */
struct node *root = NULL;        /* this points to the root */

main()
{
    char *strptr = string_space;
    int maxstrlen = MAXSTRING;
    struct node *nodeptr = node_space;
    int node_compare();
    void print_node();
    struct node **found;
    int length;

    while (fgets(strptr, maxstrlen, stdin) != NULL) {
        /* remove the trailing newline */
        length = strlen(strptr);
        strptr[length-1] = 0;
        /* set node */
        nodeptr->string = strptr;
        /* locate node into the tree */
        found = (struct node **)
            tsearch((char *) nodeptr, (char **) &root, node_compare);
        /* bump the count */
        (*found)->count++;

        if (*found == nodeptr) {
            /* node was inserted, so get a new one */
            strptr += length;
            maxstrlen -= length;
            if (maxstrlen < MINSTRING)
                break;
            if (++nodeptr >= &node_space[MAXNODES])
                break;
        }
    }
    twalk((char *)root, print_node);
}

```

```

/*
   This routine compares two nodes, based on an
   alphabetical ordering of the string field.
*/
int node_compare(node1, node2)
    struct node *node1, *node2;
{
    return strcmp(node1->string, node2->string);
}

/* Print out nodes in alphabetical order */
/*ARGSUSED2*/
void
print_node(node, order, level)
    struct node **node;
    VISIT order;
    int level;
{
    if (order == postorder || order == leaf) {
        (void) printf("string = %20s, count = %d0,
            (*node)->string, (*node)->count);
    }
}

```

**SEE ALSO**

bsearch(3), hsearch(3), lsearch(3)

**DIAGNOSTICS**

A NULL pointer is returned by `tsearch()` if there is not enough space available to create a new node.

A NULL pointer is returned by `tsearch()`, `tfind()` and `tdelete()` if `rootp` is NULL on entry.

If the datum is found, both `tsearch()` and `tfind()` return a pointer to it. If not, `tfind()` returns NULL, and `tsearch()` returns a pointer to the inserted item.

**WARNINGS**

The `root` argument to `twalk()` is one level of indirection less than the `rootp` arguments to `tsearch()` and `tdelete()`.

There are two nomenclatures used to refer to the order in which tree nodes are visited. `tsearch()` uses `preorder`, `postorder` and `endorder` to respectively refer to visiting a node before any of its children, after its left child and before its right, and after both its children. The alternate nomenclature uses `preorder`, `inorder` and `postorder` to refer to the same visits, which could result in some confusion over the meaning of `postorder`.

**BUGS**

If the calling function alters the pointer to the root, results are unpredictable.

**NAME**

`ttyname`, `isatty` – find name of a terminal

**SYNOPSIS**

```
char *ttyname(fd)
```

```
int fd;
```

```
int isatty(fd)
```

```
int fd;
```

**DESCRIPTION**

`ttyname()` returns a pointer to the null-terminated path name of the terminal device associated with file descriptor *fd*.

`isatty()` returns 1 if *fd* is associated with a terminal device, 0 otherwise.

**FILES**

`/dev/*`

**SEE ALSO**

`ctermid(3V)`, `ioctl(2)`, `ttytab(5)`

**RETURN VALUES**

On success, `ttyname()` returns a pointer to the terminal device. If *fd* does not describe a terminal device in directory `/dev`, `ttyname()` returns NULL.

`isatty()` returns 1 if *fd* is associated with a terminal device. It returns 0 otherwise.

**BUGS**

The return value points to static data which are overwritten by each call.

**NAME**

**ttyslot** – find the slot in the utmp file of the current process

**SYNOPSIS**

**int** **ttyslot**( )

**DESCRIPTION**

**ttyslot**( ) returns the index of the current user's entry in **/etc/utmp**. This is accomplished by actually scanning the file **/etc/ttytab** for the name of the terminal associated with the standard input, the standard output, or the error output (0, 1 or 2).

**RETURN VALUES**

On success, **ttyslot**( ) returns the index of the current user's entry in **/etc/utmp**. If an error was encountered while searching for the terminal name or if none of the above file descriptors is associated with a terminal device, **ttyslot**( ) returns 0.

**SYSTEM V RETURN VALUES**

If an error was encountered while searching for the terminal name or if none of the above file descriptors is associated with a terminal device, **ttyslot**( ) returns -1.

**FILES**

**/etc/ttytab**  
**/etc/utmp**



**NAME**

`ualarm` – schedule signal after interval in microseconds

**SYNOPSIS**

**unsigned ualarm(value, interval)**

**unsigned value;**

**unsigned interval;**

**DESCRIPTION**

**This is a simplified interface to `setitimer()` (see `getitimer(2)`).**

`ualarm()` sends signal `SIGALRM`, see `signal(3V)`, to the invoking process in a number of microseconds given by the *value* argument. Unless caught or ignored, the signal terminates the process.

If the *interval* argument is non-zero, the `SIGALRM` signal will be sent to the process every *interval* microseconds after the timer expires (for instance, after *value* microseconds have passed).

Because of scheduling delays, resumption of execution of when the signal is caught may be delayed an arbitrary amount. The longest specifiable delay time is 2147483647 microseconds.

The return value is the amount of time previously remaining in the alarm clock.

**SEE ALSO**

`getitimer(2)`, `sigpause(2V)`, `sigvec(2)`, `alarm(3V)`, `signal(3V)`, `sleep(3V)`, `usleep(3)`

**NAME**

ulimit – get and set user limits

**SYNOPSIS**

```
long ulimit(cmd, newlimit)
int cmd;
long newlimit;
```

**DESCRIPTION**

This function is included for System V compatibility.

This routine provides for control over process limits. The *cmd* values available are:

- 1 Get the process's file size limit. The limit is in units of 512-byte blocks and is inherited by child processes. Files of any size can be read.
- 2 Set the process's file size limit to the value of *newlimit*. Any process may decrease this limit, but only a process with an effective user ID of super-user may increase the limit. **ulimit()** will fail and the limit will be unchanged if a process with an effective user ID other than the super-user attempts to increase its file size limit.
- 3 Get the maximum possible break value. See **brk(2)**.
- 4 Get the size of the process' file descriptor table, as returned by **getdtablesize(2)**.

**RETURN VALUE**

Upon successful completion, a non-negative value is returned. Otherwise a value of -1 is returned and **errno** is set to indicate the error.

**ERRORS**

**EPERM** A user other than the super-user attempted to increase the file size limit.

**SEE ALSO**

**brk(2)**, **getdtablesize(2)**, **getrlimit(2)**, **write(2V)**

**NAME**

`ungetc` – push character back into input stream

**SYNOPSIS**

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

```
ungetc(c, stream)
```

```
FILE *stream;
```

**DESCRIPTION**

`ungetc()` pushes the character *c* back onto an input stream. That character will be returned by the next `getc()` call on that stream. `ungetc()` returns *c*, and leaves the file stream unchanged.

One character of pushback is guaranteed provided something has been read from the stream and the stream is actually buffered. In the case that stream is `stdin`, one character may be pushed back onto the buffer without a previous read statement.

If *c* equals EOF, `ungetc()` does nothing to the buffer and returns EOF.

An `fseek(3S)` erases all memory of pushed back characters.

**SEE ALSO**

`fseek(3S)`, `getc(3V)`, `setbuf(3V)`

**DIAGNOSTICS**

`ungetc()` returns EOF if it cannot push a character back.

**NAME**

usleep – suspend execution for interval in microseconds

**SYNOPSIS**

**usleep(useconds)**  
**unsigned useconds;**

**DESCRIPTION**

Suspend the current process for the number of microseconds specified by the argument. The actual suspension time may be an arbitrary amount longer because of other activity in the system, or because of the time spent in processing the call.

The routine is implemented by setting an interval timer and pausing until it occurs. The previous state of this timer is saved and restored. If the sleep time exceeds the time to the expiration of the previous timer, the process sleeps only until the signal would have occurred, and the signal is sent a short time later.

This routine is implemented using **setitimer()** (see **getitimer(2)**); it requires eight system calls each time it is invoked. A similar but less compatible function can be obtained with a single **select(2)**; it would not restart after signals, but would not interfere with other uses of **setitimer**.

**SEE ALSO**

**getitimer(2)**, **sigpause(2V)**, **alarm(3V)**, **sleep(3V)**, **ualarm(3)**

**NAME**

`utime` – set file times

**SYNOPSIS**

```
#include <utime.h>

int utime(path, times)
char *path;
struct utimbuf *times;
```

**DESCRIPTION**

`utime()` sets the access and modification times of the file named by *path*.

If *times* is NULL, the access and modification times are set to the current time. The effective user ID (UID) of the calling process must match the owner of the file or the process must have write permission for the file to use `utime()` in this manner.

If *times* is not NULL, it is assumed to point to a `utimbuf` structure, defined in `<utime.h>` as:

```
struct utimbuf {
    time_t actime; /* set the access time */
    time_t modtime; /* set the modification time */
};
```

The access time is set to the value of the first member, and the modification time is set to the value of the second member. The times contained in this structure are measured in seconds since 00:00:00 GMT Jan 1, 1970. Only the owner of the file or the super-user may use `utime()` in this manner.

Upon successful completion, `utime()` marks for update the *st\_ctime* field of the file.

**RETURN VALUES**

`utime()` returns:

- 0 on success.
- 1 on failure and sets `errno` to indicate the error.

**ERRORS**

EACCES	Search permission is denied for a component of the path prefix of <i>path</i> .
EACCES	The effective user ID is not super-user and not the owner of the file, write permission is denied for the file, and <i>times</i> is NULL.
EFAULT	<i>path</i> or <i>times</i> points outside the process's allocated address space.
EIO	An I/O error occurred while reading from or writing to the file system.
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
ENAMETOOLONG	The length of <i>path</i> exceeds {PATH_MAX}. A pathname component is longer than {NAME_MAX} while {_POSIX_NO_TRUNC} is in effect (see <code>pathconf(2V)</code> ).
ENOENT	The file referred to by <i>path</i> does not exist.
ENOTDIR	A component of the path prefix of <i>path</i> is not a directory.
EPERM	The effective user ID of the process is not super-user and not the owner of the file, and <i>times</i> is not NULL.
EROFS	The file system containing the file is mounted read-only.

**SYSTEM V ERRORS**

In addition to the above, the following may also occur:

- ENOENT *path* points to an empty string.

**SEE ALSO**

**pathconf(2V), stat(2V), utimes(2)**



**NAME**

values – machine-dependent values

**SYNOPSIS**

```
#include <values.h>
```

**DESCRIPTION**

This file contains a set of manifest constants, conditionally defined for particular processor architectures.

The model assumed for integers is binary representation (one's or two's complement), where the sign is represented by the value of the high-order bit.

<b>BITS(<i>type</i>)</b>	The number of bits in a specified type (for instance, int).
<b>HIBITS</b>	The value of a short integer with only the high-order bit set (in most implementations, 0x8000).
<b>HIBITL</b>	The value of a long integer with only the high-order bit set (in most implementations, 0x80000000).
<b>HIBITI</b>	The value of a regular integer with only the high-order bit set (usually the same as HIBITS or HIBITL).
<b>MAXSHORT</b>	The maximum value of a signed short integer (in most implementations, 0x7FFF = 32767).
<b>MAXLONG</b>	The maximum value of a signed long integer (in most implementations, 0x7FFFFFFF = 2147483647).
<b>MAXINT</b>	The maximum value of a signed regular integer (usually the same as MAXSHORT or MAXLONG).
<b>MAXFLOAT</b>	
<b>LN_MAXFLOAT</b>	The maximum value of a single-precision floating-point number, and its natural logarithm.
<b>MAXDOUBLE</b>	
<b>LN_MAXDOUBLE</b>	The maximum value of a double-precision floating-point number, and its natural logarithm.
<b>MINFLOAT</b>	
<b>LN_MINFLOAT</b>	The minimum positive value of a single-precision floating-point number, and its natural logarithm.
<b>MINDOUBLE</b>	
<b>LN_MINDOUBLE</b>	The minimum positive value of a double-precision floating-point number, and its natural logarithm.
<b>FSIGNIF</b>	The number of significant bits in the mantissa of a single-precision floating-point number.
<b>DSIGNIF</b>	The number of significant bits in the mantissa of a double-precision floating-point number.

**SEE ALSO**

intro(3), intro(3M)

**NAME**

`varargs` – handle variable argument list

**SYNOPSIS**

```
#include <varargs.h>
function(va_alist) va_dcl
va_list pvar;
va_start(pvar);
f = va_arg(pvar, type);
va_end(pvar);
```

**DESCRIPTION**

This set of macros provides a means of writing portable procedures that accept variable argument lists. Routines having variable argument lists (such as `printf(3V)`) but do not use `varargs()` are inherently nonportable, since different machines use different argument passing conventions. Routines with variable arguments lists *must* use `varargs()` functions in order to run correctly on Sun-4 systems.

`va_alist()` is used in a function header to declare a variable argument list.

`va_dcl()` is a declaration for `va_alist()`. No semicolon should follow `va_dcl()`.

`va_list()` is a type defined for the variable used to traverse the list. One such variable must always be declared.

`va_start(pvar)` is called to initialize `pvar` to the beginning of the list.

`va_arg(pvar, type)` will return the next argument in the list pointed to by `pvar`. The parameter `type` is a type name such that the type of a pointer to an object that has the specified type can be obtained simply by appending a `*` to `type`. If `type` disagrees with the type of the actual next argument (as promoted according to the default argument promotions), the behavior is undefined.

In standard C, arguments that are `char` or `short` are converted to `int` and should be accessed as `int`, arguments that are `unsigned char` or `unsigned short` are converted to `unsigned int` and should be accessed as `unsigned int`, and arguments that are `float` are converted to `double` and should be accessed as `double`. Different types can be mixed, but it is up to the routine to know what type of argument is expected, since it cannot be determined at runtime.

`va_end(pvar)` is used to finish up.

Multiple traversals, each bracketed by `va_start()` ... `va_end()`, are possible.

`va_alist()` must encompass the entire arguments list. This insures that a `#define` statement can be used to redefine or expand its value.

The argument list (or its remainder) can be passed to another function using a pointer to a variable of type `va_list()` — in which case a call to `va_arg()` in the subroutine advances the argument-list pointer with respect to the caller as well.



**EXAMPLE**

This example is a possible implementation of `execl(3V)`.

```
#include <varargs.h>
#define MAXARGS    100

/*    execl is called by
 *    execl(file, arg1, arg2, ..., (char *)0);
 */
execl (va_alist)
va_dcl
{
    va_list ap;
    char *file;
    char *args[MAXARGS];
    int argno = 0;

    va_start (ap);
    file = va_arg(ap, char *);
    while ((args[argno++] = va_arg(ap, char *)) != (char *)0)
        ;
    va_end (ap);
    return execv(file, args);
}
```

**SEE ALSO**

`execl(3V)`, `printf(3V)`

**BUGS**

It is up to the calling routine to specify how many arguments there are, since it is not possible to determine this from the stack frame. For example, `execl()` is passed a zero pointer to signal the end of the list. `printf()` can tell how many arguments are supposed to be there by the format.

The macros `va_start()` and `va_end()` may be arbitrarily complex; for example, `va_start()` might contain an opening brace, which is closed by a matching brace in `va_end()`. Thus, they should only be used where they could be placed within a single complex statement.

**NAME**

`vlimit` – control maximum system resource consumption

**SYNOPSIS**

```
#include <sys/vlimit.h>
vlimit(resource, value) int resource, value;
```

**DESCRIPTION**

**This facility is superseded by `getrlimit(2)`.**

Limits the consumption by the current process and each process it creates to not individually exceed *value* on the specified resource. If *value* is specified as `-1`, then the current limit is returned and the limit is unchanged. The resources which are currently controllable are:

<b>LIM_NORAISE</b>	A pseudo-limit; if set non-zero then the limits may not be raised. Only the super-user may remove the <i>noraise</i> restriction.
<b>LIM_CPU</b>	the maximum number of CPU-seconds to be used by each process
<b>LIM_FSIZE</b>	the largest single file which can be created
<b>LIM_DATA</b>	the maximum growth of the data+stack region using <code>sbrk()</code> (see <code>brk(2)</code> ) beyond the end of the program text
<b>LIM_STACK</b>	the maximum size of the automatically-extended stack region
<b>LIM_CORE</b>	the size of the largest core dump that will be created.
<b>LIM_MAXRSS</b>	a soft limit for the amount of physical memory (in bytes) to be given to the program. If memory is tight, the system will prefer to take memory from processes which are exceeding their declared <b>LIM_MAXRSS</b> .

Because this information is stored in the per-process information this system call must be executed directly by the shell if it is to affect all future processes created by the shell; *limit* is thus a built-in command to `cs(1)`.

The system refuses to extend the data or stack space when the limits would be exceeded in the normal way; a *break* call fails if the data space limit is reached, or the process is killed when the stack limit is reached (since the stack cannot be extended, there is no way to send a signal!).

A file I/O operation which would create a file which is too large will cause a signal `SIGXFSZ` to be generated, this normally terminates the process, but may be caught. When the cpu time limit is exceeded, a signal `SIGXCPU` is sent to the offending process; to allow it time to process the signal it is given 5 seconds grace by raising the CPU time limit.

**SEE ALSO**

`cs(1)`, `sh(1)`, `brk(2)`

**BUGS**

If **LIM\_NORAISE** is set, then no grace should be given when the CPU time limit is exceeded.

There should be *limit* and *unlimit* commands in `sh(1)` as well as in `cs(1)`.

## NAME

vprintf, vfprintf, vsprintf – print formatted output of a varargs argument list

## SYNOPSIS

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

int vprintf(format, ap)
char *format;
va_list ap;

int vfprintf(stream, format, ap)
FILE *stream;
char *format;
va_list ap;

char *vsprintf(s, format, ap)
char *s, *format;
va_list ap;
```

## SYSTEM V SYNOPSIS

```
int vsprintf(s, format, ap)
char *s, *format;
va_list ap;
```

## DESCRIPTION

vprintf(), vfprintf(), and vsprintf() are the same as printf(3V), fprintf(), and sprintf() (see printf(3V)) respectively, except that instead of being called with a variable number of arguments, they are called with an argument list as defined by varargs(3).

## RETURN VALUES

On success, vprintf() and vfprintf() return the number of characters transmitted, excluding the null character. On failure, they return EOF.

vsprintf() returns *s*.

## SYSTEM V RETURN VALUES

vsprintf() returns the number of characters transmitted, excluding the null character.

## EXAMPLES

The following demonstrates how vfprintf() could be used to write an error routine.

```
#include <stdio.h>
#include <varargs.h>
...
/* error should be called like:
 * error(function_name, format, arg1, arg2...);
 * Note: function_name and format cannot be declared
 * separately because of the definition of varargs.
 */

/*VARARGS0*/
void
error (va_alist)
    va_dcl
{
    va_list args;
    char *fmt;

    va_start(args);
    /* print name of function causing error */
```

```
(void) fprintf(stderr, "ERROR in %s: ", va_arg(args, char *));  
fmt = va_arg(args, char *);  
    /* print out remainder of message */  
(void) fprintf(stderr, fmt, args);  
va_end(args);  
(void) abort();  
}
```

SEE ALSO

printf(3V), varargs(3)

## NAME

`vsyslog` – log message with a `varargs` argument list

## SYNOPSIS

```
#include <syslog.h>
#include <varargs.h>

int vsyslog(priority, message, ap)
char *message;
va_list ap;
```

## DESCRIPTION

`vsyslog()` is the same as `syslog(3)` except that instead of being called with a variable number of arguments, it is called with an argument list as defined by `varargs(3)`.

## EXAMPLE

The following demonstrates how `vsyslog()` could be used to write an error routine.

```
#include <syslog.h>
#include <varargs.h>
...
/* error should be called like:
 *   error(pri, function_name, format, arg1, arg2...);
 * Note that pri, function_name, and format cannot be declared
 * separately because of the definition of varargs.
 */

/*VARARGS0*/
void
error(va_alist
      va_dcl;
{
    va_list args;
    int pri;
    char *message;

    va_start(args);
    pri = va_arg(args, int);
        /* log name of function causing error */
    (void) syslog(pri, "ERROR in %s", va_arg(args, char *));
    message = va_arg(args, char *);
        /* log remainder of message */
    (void) vsyslog(pri, fmt, args);
    va_end(args);
    (void) abort();
}
```

## SEE ALSO

`syslog(3)`, `varargs(3)`

## NAME

`vtimes` – get information about resource utilization

## SYNOPSIS

```
vtimes(par_vm, ch_vm)
struct vtimes *par_vm, *ch_vm;
```

## DESCRIPTION

Note: this facility is superseded by `getrusage(2)`.

`vtimes()` returns accounting information for the current process and for the terminated child processes of the current process. Either `par_vm` or `ch_vm` or both may be 0, in which case only the information for the pointers which are non-zero is returned.

After the call, each buffer contains information as defined by the contents of the include file `<sys/vtimes.h>`:

```
struct vtimes {
    int     vm_untime;      /* user time (*HZ) */
    int     vm_stime;      /* system time (*HZ) */
    /* divide next two by utime+stime to get averages */
    unsigned vm_idrss;     /* integral of d+s rss */
    unsigned vm_ixrss;     /* integral of text rss */
    int     vm_maxrss;     /* maximum rss */
    int     vm_majflt;     /* major page faults */
    int     vm_minflt;     /* minor page faults */
    int     vm_nswap;      /* number of swaps */
    int     vm_inblk;      /* block reads */
    int     vm_oublk;      /* block writes */
};
```

The `vm_untime` and `vm_stime` fields give the user and system time respectively in 60ths of a second (or 50ths if that is the frequency of wall current in your locality.) The `vm_idrss` and `vm_ixrss` measure memory usage. They are computed by integrating the number of memory pages in use each over cpu time. They are reported as though computed discretely, adding the current memory usage (in 512 byte pages) each time the clock ticks. If a process used 5 core pages over 1 cpu-second for its data and stack, then `vm_idrss` would have the value 5\*60, where `vm_untime+vm_stime` would be the 60. `vm_idrss` integrates data and stack segment usage, while `vm_ixrss` integrates text segment usage. `vm_maxrss` reports the maximum instantaneous sum of the text+data+stack core-resident page count.

The `vm_majflt` field gives the number of page faults which resulted in disk activity; the `vm_minflt` field gives the number of page faults incurred in simulation of reference bits; `vm_nswap` is the number of swaps which occurred. The number of file system input/output events are reported in `vm_inblk` and `vm_oublk`. These numbers account only for real I/O; data supplied by the caching mechanism is charged only to the first process to read or write the data.

## SEE ALSO

`getrusage(2)`, `wait(2V)`

**NAME**

xdr – library routines for external data representation

**SYNOPSIS AND DESCRIPTION**

XDR routines allow C programmers to describe arbitrary data structures in a machine-independent fashion. Data for remote procedure calls (RPC) are encoded and decoded using these routines. See [rpc\(3N\)](#).

All XDR routines require the header `<rpc/xdr.h>` to be included.

The XDR routines have been grouped by usage on the following man pages.

**xdr\_admin(3N)** Library routines for managing the XDR stream. The routines documented on this page include:

```
xdr_getpos()
xdr_inline()
xdrrec_endofrecord()
xdrrec_eof()
xdrrec_readbytes()
xdrrec_skiprecord()
xdr_setpos()
```

**xdr\_complex(3N)** Library routines for translating complex data types into their external data representation. The routines documented on this page include:

```
xdr_array()
xdr_bytes()
xdr_opaque()
xdr_pointer()
xdr_reference()
xdr_string()
xdr_union()
xdr_vector()
xdr_wrapstring()
```

**xdr\_create(3N)** Library routines for creating XDR streams. The routines documented on this page include:

```
xdr_destroy()
xdrmem_create()
xdrrec_create()
xdrstdio_create()
```

**xdr\_simple(3N)** Library routines for translating simple data types into their external data representation. The routines documented on this page include:

```
xdr_bool()
xdr_char()
xdr_double()
xdr_enum()
xdr_float()
xdr_free()
xdr_int()
xdr_long()
xdr_short()
xdr_u_char()
xdr_u_int()
xdr_u_long()
xdr_u_short()
xdr_void()
```

**SEE ALSO**

**rpc(3N), xdr\_admin(3N), xdr\_complex(3N), xdr\_create(3N), xdr\_simple(3N)**

*Network Programming*



## NAME

xdr\_getpos, xdr\_inline, xdrrec\_endofrecord, xdrrec\_eof, xdrrec\_readbytes, xdrrec\_skiprecord, xdr\_setpos  
 – library routines for management of the XDR stream

## DESCRIPTION

XDR library routines allow C programmers to describe arbitrary data structures in a machine-independent fashion. Protocols such as remote procedure calls (RPC) use these routines to describe the format of the data.

These routines deal specifically with the management of the XDR stream.

## Routines

The XDR data structure is defined in the RPC/XDR Library Definitions of the *Network Programming*.

```
#include <rpc/xdr.h>
```

```
u_int xdr_getpos(xdrs)
```

```
XDR *xdrs;
```

Invoke the get-position routine associated with the XDR stream, *xdrs*. The routine returns an unsigned integer, which indicates the position of the XDR byte stream. A desirable feature of XDR streams is that simple arithmetic works with this number, although the XDR stream instances need not guarantee this.

```
long * xdr_inline(xdrs, len)
```

```
XDR *xdrs;
```

```
int len;
```

Invoke the in-line routine associated with the XDR stream, *xdrs*. The routine returns a pointer to a contiguous piece of the stream's buffer; *len* is the byte length of the desired buffer. Note: A pointer is cast to long \*.

Warning: *xdr\_inline()* may return NULL if it cannot allocate a contiguous piece of a buffer. Therefore the behavior may vary among stream instances; it exists for the sake of efficiency.

```
bool_t xdrrec_endofrecord(xdrs, sendnow)
```

```
XDR *xdrs;
```

```
int sendnow;
```

This routine can be invoked only on streams created by *xdrrec\_create()* (see *xdr\_create(3N)*). The data in the output buffer is marked as a completed record, and the output buffer is optionally written out if *sendnow* is non-zero. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdrrec_eof(xdrs)
```

```
XDR *xdrs;
```

```
int empty;
```

This routine can be invoked only on streams created by *xdrrec\_create()* (see *xdr\_create(3N)*). After consuming the rest of the current record in the stream, this routine returns TRUE if the stream has no more input, FALSE otherwise.

```
int xdrrec_readbytes(xdrs, addr, nbytes)
```

```
XDR *xdrs;
```

```
caddr_t addr;
```

```
u_int nbytes;
```

This routine can be invoked only on streams created by *xdrrec\_create()* (see *xdr\_create(3N)*). It attempts to read *nbytes* bytes from the XDR stream into the buffer pointed to by *addr*. On success it returns the number of bytes read. Returns -1 on failure. A return value of 0 indicates an end of record.

**bool\_t xdrrec\_skiprecord(xdrs)****XDR \*xdrs;**

This routine can be invoked only on streams created by **xdrrec\_create()** (see **xdr\_create(3N)**). It tells the XDR implementation that the rest of the current record in the stream's input buffer should be discarded. This routine returns TRUE if it succeeds, FALSE otherwise.

**bool\_t xdr\_setpos(xdrs, pos)****XDR \*xdrs;****u\_int pos;**

Invoke the set position routine associated with the XDR stream *xdrs*. The parameter *pos* is a position value obtained from **xdr\_getpos()**. This routine returns 1 if the XDR stream could be repositioned, and 0 otherwise.

Warning: It is difficult to reposition some types of XDR streams, so this routine may fail with one type of stream and succeed with another.

**SEE ALSO****xdr(3N), xdr\_complex(3N), xdr\_create(3N), xdr\_simple(3N)**

**NAME**

xdr\_array, xdr\_bytes, xdr\_opaque, xdr\_pointer, xdr\_reference, xdr\_string, xdr\_union, xdr\_vector, xdr\_wrapstring – library routines for translating complex data types

**DESCRIPTION**

XDR library routines allow C programmers to describe complex data structures in a machine-independent fashion. Protocols such as remote procedure calls (RPC) use these routines to describe the format of the data.

**Routines**

The XDR data structure is defined in the RPC/XDR Library Definitions of the *Network Programming*.

```
#include <rpc/xdr.h>
```

```
bool_t xdr_array(xdrs, arrp, sizep, maxsize, elsize, elproc)
```

```
XDR *xdrs;
```

```
char **arrp;
```

```
u_int *sizep, maxsize, elsize;
```

```
xdrproc_t elproc;
```

A filter primitive that translates between a variable-length array and its corresponding external representations. The parameter *arrp* is the address of the pointer to the array, while *sizep* is the address of the element count of the array. This value is used by the filter while encoding and is set by it while decoding; the routine fails if the element count exceeds *maxsize*. The parameter *elsize* is the *sizeof* each of the array's elements, and *elproc* is an XDR filter that translates between the array elements' C form, and their external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_bytes(xdrs, arrp, sizep, maxsize)
```

```
XDR *xdrs;
```

```
char **arrp;
```

```
u_int *sizep, maxsize;
```

A filter primitive that translates between an array of bytes and its external representation. It treats the array of bytes as opaque data. The parameter *arrp* is the address of the array of bytes. While decoding if *\*arrp* is NULL, then the necessary storage is allocated to hold the array. This storage can be freed by using `xdr_free()` (see `xdr_simple(3N)`). *sizep* is the pointer to the actual length specifier for the array. This value is used by the filter while encoding and is set by it when decoding. *maxsize* is the maximum length of the array. The routine fails if the actual length of the array is greater than *maxsize*. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_opaque(xdrs, cp, cnt)
```

```
XDR *xdrs;
```

```
char *cp;
```

```
u_int cnt;
```

A filter primitive that translates between fixed size opaque data and its external representation. The parameter *cp* is the address of the opaque object, and *cnt* is its size in bytes. This routine returns TRUE if it succeeds, FALSE otherwise.

**bool\_t xdr\_pointer(xdrs, objpp, objsize, objproc)**

XDR \*xdrs;  
char \*\*objpp;  
u\_int objsize;  
xdrproc\_t objproc;

Like `xdr_reference()` except that it serializes NULL pointers, whereas `xdr_reference()` does not. Thus, `xdr_pointer()` can represent recursive data structures, such as binary trees or linked lists. The parameter `objpp` is the address of the pointer; `objsize` is the *sizeof* the structure that `*objpp` points to; and `objproc` is an XDR procedure that filters the structure between its C form and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

**bool\_t xdr\_reference(xdrs, pp, size, proc)**

XDR \*xdrs;  
char \*\*pp;  
u\_int size;  
xdrproc\_t proc;

A primitive that provides pointer chasing within structures. The parameter `pp` is the address of the pointer; `size` is the *sizeof* the structure that `*pp` points to; and `proc` is an XDR procedure that filters the structure between its C form and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

Warning: This routine does not understand NULL pointers. Use `xdr_pointer()` instead.

**bool\_t xdr\_string(xdrs, strp, maxsize)**

XDR \*xdrs;  
char \*\*strp;  
u\_int maxsize;

A filter primitive that translates between C strings and their corresponding external representations. The routine fails if the string being translated is longer than `maxsize`. `strp` is the address of the pointer to the string. While decoding if `*strp` is NULL, then the necessary storage is allocated to hold this null-terminated string and `*strp` is set to point to this. This storage can be freed by using `xdr_free()` (see `xdr_simple(3N)`). This routine returns TRUE if it succeeds, FALSE otherwise.

**bool\_t xdr\_union(xdrs, dscmp, unp, choices, defaultarm)**

XDR \*xdrs;  
int \*dscmp;  
char \*unp;  
struct xdr\_discrim \*choices;  
bool\_t (\*defaultarm) (); /\* may be NULL \*/

A filter primitive that translates between a discriminated C union and its corresponding external representation. It first translates the discriminant of the union located at `dscmp`. This discriminant is always an `enum_t`. Next the union located at `unp` is translated. The parameter `choices` is a pointer to an array of `xdr_discrim` structures. Each structure contains an ordered pair of [`value,proc`]. If the union's discriminant is equal to any of the `values`, then the associated `proc` is called to translate the union. The end of the `xdr_discrim` structure array is denoted by a NULL pointer. If the discriminant is not found in the `choices` array, then the `defaultarm` procedure is called (if it is not NULL). This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_vector(xdrs, arrp, size, elsize, elproc)
XDR *xdrs;
char *arrp;
u_int size, elsize;
xdrproc_t elproc;
```

A filter primitive that translates between fixed-length arrays and their corresponding external representations. The parameter *arrp* is the address of the array, while *size* is the element count of the array. The parameter *elsize* is the *sizeof* each of the array's elements, and *elproc* is an XDR filter that translates between the array elements' C form, and their external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_wrapstring(xdrs, strp)
XDR *xdrs;
char **strp;
```

A primitive that calls `xdr_string(xdrs, strp, MAXUNSIGNED)`; where `MAXUNSIGNED` is the maximum value of an unsigned integer. `xdr_wrapstring()` is handy because the RPC package passes a maximum of two XDR routines as parameters, and `xdr_string()`, one of the most frequently used primitives, requires three. *strp* is the address of the pointer to the string. While decoding if *\*strp* is NULL, then the necessary storage is allocated to hold the null-terminated string and *\*strp* is set to point to this. This storage can be freed by using `xdr_free()` (see `xdr_simple(3N)`). This routine returns TRUE if it succeeds, FALSE otherwise.

SEE ALSO

`xdr(3N)`, `xdr_admin(3N)`, `xdr_create(3N)`, `xdr_simple(3N)`

**NAME**

`xdr_destroy`, `xdrmem_create`, `xdrrec_create`, `xdrstdio_create` – library routines for external data representation stream creation

**DESCRIPTION**

XDR library routines allow C programmers to describe arbitrary data structures in a machine-independent fashion. Protocols such as remote procedure calls (RPC) use these routines to describe the format of the data.

These routines deal with the creation of XDR streams. XDR streams have to be created before any data can be translated into XDR format.

**Routines**

The `XDR`, `CLIENT`, and `SVCXPRT` data structures are defined in the RPC/XDR Library Definitions of the *Network Programming*.

```
#include <rpc/xdr.h>
```

```
void xdr_destroy(xdrs)
XDR *xdrs;
```

Invoke the destroy routine associated with the XDR stream, `xdrs`. Destruction usually involves freeing private data structures associated with the stream. Using `xdrs` after invoking `xdr_destroy()` is undefined.

```
void xdrmem_create(xdrs, addr, size, op)
XDR *xdrs;
char *addr;
u_int size;
enum xdr_op op;
```

This routine initializes the XDR stream object pointed to by `xdrs`. The stream's data is written to, or read from, a chunk of memory at location `addr` whose length is no more than `size` bytes long. `size` should be a multiple of 4. The `op` determines the direction of the XDR stream (either `XDR_ENCODE`, `XDR_DECODE`, or `XDR_FREE`).

```
void xdrrec_create(xdrs, sendsz, recvsz, handle, readit, writeit)
XDR *xdrs;
u_int sendsz, recvsz;
char *handle;
int (*readit) (), (*writeit) ();
```

This routine initializes the XDR stream object pointed to by `xdrs`. The stream's data is written to a buffer of size `sendsz`; a value of zero indicates the system should use a suitable default. The stream's data is read from a buffer of size `recvsz`; it too can be set to a suitable default by passing a zero value. When a stream's output buffer is full, `writeit` is called. Similarly, when a stream's input buffer is empty, `readit` is called. The behavior of these two routines is similar to `read(2V)` and `write(2V)`, except that `handle` is passed to the former routines as the first parameter. Note: The XDR stream's `op` field must be set by the caller. `sendsz` and `recvsz` should be multiples of 4.

Warning: This XDR stream implements an intermediate record stream. Therefore there are additional bytes in the stream to provide record boundary information.

```
void xdrstdio_create(xdrs, filep, op)
XDR *xdrs;
FILE *filep;
enum xdr_op op;
```

This routine initializes the XDR stream object pointed to by *xdrs*. The XDR stream data is written to, or read from, the Standard I/O stream *filep*. The parameter *op* determines the direction of the XDR stream (either **XDR\_ENCODE**, **XDR\_DECODE**, or **XDR\_FREE**).

Warning: The destroy routine associated with such XDR streams calls **fflush()** on the *file* stream, but never **fclose(3V)**.

**SEE ALSO**

**read(2V)**, **write(2V)**, **fclose(3V)**, **xdr(3N)**, **xdr\_admin(3N)**, **xdr\_complex(3N)**, **xdr\_simple(3N)**

**NAME**

xdr\_bool, xdr\_char, xdr\_double, xdr\_enum, xdr\_float, xdr\_free, xdr\_int, xdr\_long, xdr\_short, xdr\_u\_char, xdr\_u\_int, xdr\_u\_long, xdr\_u\_short, xdr\_void – library routines for translating simple data types

**DESCRIPTION**

XDR library routines allow C programmers to describe simple data structures in a machine-independent fashion. Protocols such as remote procedure calls (RPC) use these routines to describe the format of the data.

These routines require the creation of XDR streams (see `xdr_create(3N)`).

**Routines**

The XDR data structure is defined in the RPC/XDR Library Definitions of the *Network Programming*.

```
#include <rpc/xdr.h>
```

```
bool_t xdr_bool(xdrs, bp)
XDR *xdrs;
bool_t *bp;
```

A filter primitive that translates between a boolean (C integer) and its external representation. When encoding data, this filter produces values of either one or zero. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_char(xdrs, cp)
XDR *xdrs;
char *cp;
```

A filter primitive that translates between a C character and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

Note: Encoded characters are not packed, and occupy 4 bytes each. For arrays of characters, it is worthwhile to consider `xdr_bytes()`, `xdr_opaque()` or `xdr_string()`, see `xdr_complex(3N)`.

```
bool_t xdr_double(xdrs, dp)
XDR *xdrs;
double *dp;
```

A filter primitive that translates between a C **double** precision number and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_enum(xdrs, ep)
XDR *xdrs;
enum_t *ep;
```

A filter primitive that translates between a C **enum** (actually integer) and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_float(xdrs, fp)
XDR *xdrs;
float *fp;
```

A filter primitive that translates between a C **float** and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.



```
void xdr_free(proc, objp)
xdrproc_t proc;
char *objp;
```

Generic freeing routine. The first argument is the XDR routine for the object being freed. The second argument is a pointer to the object itself. Note: The pointer passed to this routine is *not* freed, but what it points to *is* freed, recursively such that objects pointed to are also freed for example, linked lists.

```
bool_t xdr_int(xdrs, ip)
XDR *xdrs;
int *ip;
```

A filter primitive that translates between a C **integer** and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_long(xdrs, lp)
XDR *xdrs;
long *lp;
```

A filter primitive that translates between a C **long** integer and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_short(xdrs, sp)
XDR *xdrs;
short *sp;
```

A filter primitive that translates between a C **short** integer and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_u_char(xdrs, ucp)
XDR *xdrs;
unsigned char *ucp;
```

A filter primitive that translates between an **unsigned** C character and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_u_int(xdrs, up)
XDR *xdrs;
unsigned *up;
```

A filter primitive that translates between a C **unsigned** integer and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_u_long(xdrs, ulp)
XDR *xdrs;
unsigned long *ulp;
```

A filter primitive that translates between a C **unsigned long** integer and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_u_short(xdrs, usp)
XDR *xdrs;
unsigned short *usp;
```

A filter primitive that translates between a C **unsigned short** integer and its external representation. This routine returns TRUE if it succeeds, FALSE otherwise.

```
bool_t xdr_void()
```

This routine always returns TRUE. It may be passed to RPC routines that require a function parameter, where nothing is to be done.

**SEE ALSO**

**xdr(3N), xdr\_admin(3N), xdr\_complex(3N), xdr\_create(3N)**

## NAME

ypclnt, yp\_get\_default\_domain, yp\_bind, yp\_unbind, yp\_match, yp\_first, yp\_next, yp\_all, yp\_order, yp\_master, yperr\_string, ypprot\_err – NIS client interface

## SYNOPSIS AND DESCRIPTION

This package of functions provides an interface to the Network Information Service (NIS). The package can be loaded from the standard library, `/usr/lib/libc.a`. Refer to `ypfiles(5)` and `ypserv(8)` for an overview of the NIS name service, including the definitions of *map* and *domain*, and a description of the various servers, databases, and commands that comprise the NIS services.

All input parameters names begin with *in*. Output parameters begin with *out*. Output parameters of type `char **` should be addresses of uninitialized character pointers. Memory is allocated by the NIS client package using `malloc(3V)`, and may be freed if the user code has no continuing need for it. For each *outkey* and *outval*, two extra bytes of memory are allocated at the end that contain NEWLINE and the null character, respectively, but these two bytes are not reflected in *outkeylen* or *outvallen*. *indomain* and *inmap* strings must not be empty and must be null-terminated. String parameters which are accompanied by a count parameter may not be NULL, but may point to null strings, with the count parameter indicating this. Counted strings need not be null-terminated.

All functions in this package of type *int* return 0 if they succeed, and a failure code (`YPERR_xxxx`) otherwise. Failure codes are described under DIAGNOSTICS below.

**yp\_bind (indomain);**  
**char \*indomain;**

To use the NIS services, the client process must be “bound” to a NIS server that serves the appropriate domain using `yp_bind()`. Binding need not be done explicitly by user code; this is done automatically whenever a NIS lookup function is called. `yp_bind()` can be called directly for processes that make use of a backup strategy (for example, a local file) in cases when NIS services are not available.

**void**  
**yp\_unbind (indomain)**  
**char \*indomain;**

Each binding allocates (uses up) one client process socket descriptor; each bound domain costs one socket descriptor. However, multiple requests to the same domain use that same descriptor. `yp_unbind()` is available at the client interface for processes that explicitly manage their socket descriptors while accessing multiple domains. The call to `yp_unbind()` make the domain *unbound*, and free all per-process and per-node resources used to bind it.

If an RPC failure results upon use of a binding, that domain will be unbound automatically. At that point, the `ypclnt` layer will retry forever or until the operation succeeds, provided that `ypbind` is running, and either

- a) the client process cannot bind a server for the proper domain, or
- b) RPC requests to the server fail.

If an error is not RPC-related, or if `ypbind` is not running, or if a bound `ypserv` process returns any answer (success or failure), the `ypclnt` layer will return control to the user code, either with an error code, or a success code and any results.

```
yp_get_default_domain (outdomain);  
char **outdomain;
```

The NIS lookup calls require a map name and a domain name, at minimum. It is assumed that the client process knows the name of the map of interest. Client processes should fetch the node's default domain by calling `yp_get_default_domain()`, and use the returned *outdomain* as the *indomain* parameter to successive NIS calls.

```
yp_match(indomain, inmap, inkey, inkeylen, outval, outvallen)  
char *indomain;  
char *inmap;  
char *inkey;  
int inkeylen;  
char **outval;  
int *outvallen;
```

`yp_match()` returns the value associated with a passed key. This key must be exact; no pattern matching is available.

```
yp_first(indomain, inmap, outkey, outkeylen, outval, outvallen)  
char *indomain;  
char *inmap;  
char **outkey;  
int *outkeylen;  
char **outval;  
int *outvallen;
```

`yp_first()` returns the first key-value pair from the named map in the named domain.

```
yp_next(indomain, inmap, inkey, inkeylen, outkey, outkeylen, outval, outvallen);  
char *indomain;  
char *inmap;  
char *inkey;  
int inkeylen;  
char **outkey;  
int *outkeylen;  
char **outval;  
int *outvallen;
```

`yp_next()` returns the next key-value pair in a named map. The *inkey* parameter should be the *outkey* returned from an initial call to `yp_first()` (to get the second key-value pair) or the one returned from the *n*th call to `yp_next()` (to get the *n*th + second key-value pair).

The concept of first (and, for that matter, of next) is particular to the structure of the NIS map being processed; there is no relation in retrieval order to either the lexical order within any original (non-NIS) data base, or to any obvious numerical sorting order on the keys, values, or key-value pairs. The only ordering guarantee made is that if the `yp_first()` function is called on a particular map, and then the `yp_next()` function is repeatedly called on the same map at the same server until the call fails with a reason of `YPERR_NOMORE`, every entry in the data base will be seen exactly once. Further, if the same sequence of operations is performed on the same map at the same server, the entries will be seen in the same order.

Under conditions of heavy server load or server failure, it is possible for the domain to become unbound, then bound once again (perhaps to a different server) while a client is running. This can cause a break in one of the enumeration rules; specific entries may be seen twice by the client, or not at all. This approach protects the client from error messages that would otherwise be returned in the midst of the enumeration. The next paragraph describes a better solution to enumerating all entries in a map.

```
yp_all(indomain, inmap, incallback);
char *indomain;
char *inmap;
struct ypall_callback *incallback;
```

`yp_all()` provides a way to transfer an entire map from server to client in a single request using TCP (rather than UDP as with other functions in this package). The entire transaction take place as a single RPC request and response. You can use `yp_all()` just like any other NIS procedure, identify the map in the normal manner, and supply the name of a function which will be called to process each key-value pair within the map. You return from the call to `yp_all()` only when the transaction is completed (successfully or unsuccessfully), or your `foreach` function decides that it does not want to see any more key-value pairs.

The third parameter to `yp_all()` is

```
struct ypall_callback *incallback {
    int (*foreach)();
    char *data;
};
```

The function `foreach` is called

```
foreach(instatus, inkey, inkeylen, inval, invallen, indata);
int instatus;
char *inkey;
int inkeylen;
char *inval;
int invallen;
char *indata;
```

The `instatus` parameter will hold one of the return status values defined in `<rpcsvc/yp_prot.h>` — either `YP_TRUE` or an error code. See `ypprot_err()`, below, for a function which converts a NIS protocol error code to a `ypclnt` layer error code.

The key and value parameters are somewhat different than defined in the synopsis section above. First, the memory pointed to by the `inkey` and `inval` parameters is private to the `yp_all()` function, and is overwritten with the arrival of each new key-value pair. It is the responsibility of the `foreach` function to do something useful with the contents of that memory, but it does not own the memory itself. Key and value objects presented to the `foreach` function look exactly as they do in the server's map — if they were not NEWLINE-terminated or null-terminated in the map, they will not be here either.

The `indata` parameter is the contents of the `incallback->data` element passed to `yp_all()`. The `data` element of the callback structure may be used to share state information between the `foreach` function and the mainline code. Its use is optional, and no part of the NIS client package inspects its contents — cast it to something useful, or ignore it as you see fit.

The `foreach` function is a Boolean. It should return zero to indicate that it wants to be called again for further received key-value pairs, or non-zero to stop the flow of key-value pairs. If `foreach` returns a non-zero value, it is not called again; the functional value of `yp_all()` is then 0.

```

yp_order(indomain, inmap, outorder);
char *indomain;
char *inmap;
int *outorder;

```

yp\_order() returns the order number for a map.

```

yp_master(indomain, inmap, outname);
char *indomain;
char *inmap;
char **outname;

```

yp\_master() returns the machine name of the master NIS server for a map.

```

char *yperr_string(incode)
int incode;

```

yperr\_string() returns a pointer to an error message string that is null-terminated but contains no period or NEWLINE.

```

ypprot_err(incode)
unsigned int incode;

```

ypprot\_err() takes a NIS protocol error code as input, and returns a ypclnt layer error code, which may be used in turn as an input to yperr\_string().

#### FILES

```

<rpcsvc/ypclnt.h>
<rpcsvc/yp_prot.h>
/usr/lib/libc.a

```

#### SEE ALSO

malloc(3V), ypupdate(3N), ypfiles(5), ypserv(8)

#### DIAGNOSTICS

All integer functions return 0 if the requested operation is successful, or one of the following errors if the operation fails.

```

#define YPERR_BADARGS
    1      /* args to function are bad */

#define YPERR_RPC
    2      /* RPC failure - domain has been unbound */

#define YPERR_DOMAIN
    3      /* can't bind to server on this domain */

#define YPERR_MAP
    4      /* no such map in server's domain */

#define YPERR_KEY
    5      /* no such key in map */

#define YPERR_YPERR
    6      /* internal yp server or client error */

#define YPERR_RESRC
    7      /* resource allocation failure */

#define YPERR_NOMORE
    8      /* no more records in map database */

#define YPERR_PMAP
    9      /* can't communicate with portmapper */

#define YPERR_YPBIND

```

```
    10    /* can't communicate with ypbind */  
#define YPERR_YPSErv  
    11    /* can't communicate with ypserv */  
#define YPERR_NODOM  
    12    /* local domain name not set */  
#define YPERR_BADDBfR  
    13    /* yp database is bad */  
#define YPERR_VERSfR  
    14    /* yp version mismatch */  
#define YPERR_ACCESS  
    15    /* access violation */  
#define YPERR_BUSY  
    16    /* database busy */
```

**NOTES**

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications plc, and may not be used without permission.

**NAME**

`yp_update` – changes NIS information

**SYNOPSIS**

```
#include <rpcsvc/ypclnt.h>

yp_update(domain, map, ypop, key, keylen, data, datalen)
char *domain;
char *map;
unsigned ypop
char *key;
int keylen;
char *data;
int datalen;
```

**DESCRIPTION**

`yp_update()` is used to make changes to the Network Information Service (NIS) database. The syntax is the same as that of `yp_match()` (see `ypclnt(3N)`) except for the extra parameter `ypop` which may take on one of four values. If it is `YPOP_CHANGE` then the data associated with the key will be changed to the new value. If the key is not found in the database, then `yp_update()` returns `YPERR_KEY`. If `ypop` has the value `YPOP_INSERT` then the key-value pair will be inserted into the database. The error `YPERR_KEY` is returned if the key already exists in the database. To store an item into the database without concern for whether it exists already or not, pass `ypop` as `YPOP_STORE` and no error will be returned if the key already or does not exist. To delete an entry, the value of `ypop` should be `YPOP_DELETE`.

This routine depends upon secure RPC, and will not work unless the network is running secure RPC.

**SEE ALSO**

`ypclnt(3N)`

*System and Network Administration*

**NOTES**

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications plc, and may not be used without permission.



## NAME

intro – introduction to the lightweight process library (LWP)

## DESCRIPTION

The lightweight process library (LWP) provides a mechanism to support multiple threads of control that share a single address space. Under SunOS, the address space is derived from a single *forked* (“heavy-weight”) process. Each thread has its own stack segment (specified when the thread is created) so that it can access local variables and make procedure calls independently of other threads. The collection of threads sharing an address space is called a *pod*. Under SunOS, threads share all of the resources of the heavyweight process that contains the pod, including descriptors and signal handlers.

The LWP provides a means for creating and destroying threads, message exchange between threads, manipulating condition variables and monitors, handling synchronous exceptions, mapping asynchronous events into messages, mapping synchronous events into exceptions, arranging for special per-thread context, multiplexing the clock for timeouts, and scheduling threads both preemptively and non-preemptively.

The LWP system exists as a library of routines (`/usr/lib/liblwp.a`) linked in (`-llwp`) with a client program which should `#include` the file `<lwp/lwp.h>`. `main` is transparently converted into a lightweight process as soon as it attempts to use any LWP primitives.

When an object created by a LWP primitive is destroyed, every attempt is made to clean up after it. For example, if a thread dies, all threads blocked on sends to or receives from that thread are unblocked, and all monitor locks held by the dead thread are released.

Because there is no kernel support for threads at present, system calls effectively block the entire pod. By linking in the non-blocking I/O library (`-lnbio`) ahead of the LWP library, you can alleviate this problem for those system calls that can issue a signal when a system call would be profitable to try. This library (which redefines some system calls) uses asynchronous I/O and events (for example, `SIGCHLD` and `SIGIO`) to make blocking less painful. The system calls remapped by the `nbio` library are: `open(2V)`, `socket(2)`, `pipe(2V)`, `close(2V)`, `read(2V)`, `write(2V)`, `send(2)`, `recv(2)`, `accept(2)`, `connect(2)`, `select(2)` and `wait(2V)`.

## RETURN VALUES

LWP primitives return non-negative integers on success. On errors, they return `-1`. See `lwp_perror(3L)` for details on error handling.

## FILES

`/usr/lib/liblwp.a`  
`/usr/lib/libnbio.a`

## SEE ALSO

`accept(2)`, `close(2V)`, `connect(2)`, `open(2V)`, `pipe(2V)`, `read(2V)`, `recv(2)`, `select(2)`, `send(2)`, `socket(2)`, `wait(2V)` `write(2V)`  
*Lightweight Processes in the System Services Overview*

## INDEX

The following are the primitives currently supported, grouped roughly by function.

## Thread Creation

`lwp_self(tid)`  
`lwp_getstate(tid, statvec)`  
`lwp_setregs(tid, machstate)`  
`lwp_getregs(tid, machstate)`  
`lwp_ping(tid)`  
`lwp_create(tid, pc, prio, flags, stack, nargs, arg1, ..., argn)`  
`lwp_destroy(tid)`  
`lwp_enumerate(vec, maxsize)`  
`pod_setexit(status)`  
`pod_getexit()`  
`pod_exit(status)`  
`SAMETHREAD(t1, t2)`

**Thread Scheduling**

**pod\_setmaxpri(maxprio)**  
**pod\_getmaxpri()**  
**pod\_getmaxsize()**  
**lwp\_resched(prio)**  
**lwp\_setpri(tid, prio)**  
**lwp\_sleep(timeout)**  
**lwp\_suspend(tid)**  
**lwp\_resume(tid)**  
**lwp\_yield(tid)**  
**lwp\_join(tid)**

**Error Handling**

**lwp\_geterr()**  
**lwp\_perror(s)**  
**lwp\_errstr()**

**Messages**

**msg\_send(tid, argbuf, argsize, resbuf, ressize)**  
**msg\_recv(tid, argbuf, argsize, resbuf, ressize, timeout)**  
**MSG\_RECVALL(tid, argbuf, argsize, resbuf, ressize, timeout)**  
**msg\_reply(tid)**  
**msg\_enumsend(vec, maxsize)**  
**msg\_enumrecv(vec, maxsize)**

**Event Mapping (Agents)**

**agt\_create(agt, event, memory)**  
**agt\_enumerate(vec, maxsize)**  
**agt\_trap(event)**

**Thread Synchronization: Monitors**

**mon\_create(mid)**  
**mon\_destroy(mid)**  
**mon\_enter(mid)**  
**mon\_exit(mid)**  
**mon\_enumerate(vec, maxsize)**  
**mon\_waiters (mid, owner, vec, maxsize)**  
**mon\_cond\_enter(mid)**  
**mon\_break(mid)**  
**MONITOR(mid)**  
**SAMEMON(m1, m2)**

**Thread Synchronization: Condition Variables**

**cv\_create(cv, mid)**  
**cv\_destroy(cv)**  
**cv\_wait(cv)**  
**cv\_notify(cv)**  
**cv\_send(cv, tid)**  
**cv\_broadcast(cv)**  
**cv\_enumerate(vec, maxsize)**  
**cv\_waiters(cv, vec, maxsize)**  
**SAMECV(c1, c2)**

**Exception Handling**

**exc\_handle(pattern, func, arg)**  
**exc\_unhandle()**  
**(\*exc\_bound(pattern, arg))()**  
**exc\_notify(pattern)**  
**exc\_raise(pattern)**

`exc_on_exit(func, arg)`

`exc_uniqpatt()`

#### Special Context Handling

`lwp_ctxinit(tid, cookie)`

`lwp_ctxremove(tid, cookie)`

`lwp_ctxset(save, restore, ctxsize, optimise)`

`lwp_ctxmemget(mem, tid, ctx)`

`lwp_ctxmemset(mem, tid, ctx)`

`lwp_fpset(tid)`

`lwp_libcset(tid)`

#### Stack Management

`CHECK(location, result)`

`lwp_setstkcache(minsize, numstks)`

`lwp_newstk()`

`lwp_datastk(data, size, addr)`

`lwp_stkcswwset(tid, limit)`

`lwp_checkstkset(tid, limit)`

`STKTOP(s)`

#### BUGS

There is no language support available from C.

There is no kernel support yet. Thus system calls in different threads cannot execute in parallel.

Killing a process that uses the non-blocking I/O library may leave objects (such as its standard input) in a non-blocking state. This could cause confusion to the shell.

#### LIST OF LWP LIBRARY FUNCTIONS

Name	Appears on Page	Description
<code>agt_create</code>	<code>agt_create(3L)</code>	map LWP events into messages
<code>agt_enumerate</code>	<code>agt_create(3L)</code>	map LWP events into messages
<code>agt_trap</code>	<code>agt_create(3L)</code>	map LWP events into messages
<code>CHECK</code>	<code>lwp_newstk(3L)</code>	LWP stack management
<code>cv_broadcast</code>	<code>cv_create(3L)</code>	manage LWP condition variables
<code>cv_create</code>	<code>cv_create(3L)</code>	manage LWP condition variables
<code>cv_destroy</code>	<code>cv_create(3L)</code>	manage LWP condition variables
<code>cv_enumerate</code>	<code>cv_create(3L)</code>	manage LWP condition variables
<code>cv_notify</code>	<code>cv_create(3L)</code>	manage LWP condition variables
<code>cv_send</code>	<code>cv_create(3L)</code>	manage LWP condition variables
<code>cv_wait</code>	<code>cv_create(3L)</code>	manage LWP condition variables
<code>cv_waiters</code>	<code>cv_create(3L)</code>	manage LWP condition variables
<code>exc_bound</code>	<code>exc_handle(3L)</code>	LWP exception handling
<code>exc_handle</code>	<code>exc_handle(3L)</code>	LWP exception handling
<code>exc_notify</code>	<code>exc_handle(3L)</code>	LWP exception handling
<code>exc_on_exit</code>	<code>exc_handle(3L)</code>	LWP exception handling
<code>exc_raise</code>	<code>exc_handle(3L)</code>	LWP exception handling
<code>exc_unhandle</code>	<code>exc_handle(3L)</code>	LWP exception handling
<code>exc_uniqpatt</code>	<code>exc_handle(3L)</code>	LWP exception handling
<code>lwp_checkstkset</code>	<code>lwp_newstk(3L)</code>	LWP stack management
<code>lwp_create</code>	<code>lwp_create(3L)</code>	LWP thread creation and destruction primitives
<code>lwp_ctxinit</code>	<code>lwp_ctxinit(3L)</code>	special LWP context operations
<code>lwp_ctxmemget</code>	<code>lwp_ctxinit(3L)</code>	special LWP context operations
<code>lwp_ctxmemset</code>	<code>lwp_ctxinit(3L)</code>	special LWP context operations
<code>lwp_ctxremove</code>	<code>lwp_ctxinit(3L)</code>	special LWP context operations
<code>lwp_ctxset</code>	<code>lwp_ctxinit(3L)</code>	special LWP context operations
<code>lwp_datastk</code>	<code>lwp_newstk(3L)</code>	LWP stack management

<b>lwp_destroy</b>	<b>lwp_create(3L)</b>	LWP thread creation and destruction primitives
<b>lwp_enumerate</b>	<b>lwp_status(3L)</b>	LWP status information
<b>lwp_errstr</b>	<b>lwp_perror(3L)</b>	LWP error handling
<b>lwp_fpset</b>	<b>lwp_ctxinit(3L)</b>	special LWP context operations
<b>lwp_geterr</b>	<b>lwp_perror(3L)</b>	LWP error handling
<b>lwp_getregs</b>	<b>lwp_status(3L)</b>	LWP status information
<b>lwp_getstate</b>	<b>lwp_status(3L)</b>	LWP status information
<b>lwp_join</b>	<b>lwp_yield(3L)</b>	control LWP scheduling
<b>lwp_libcset</b>	<b>lwp_ctxinit(3L)</b>	special LWP context operations
<b>lwp_newstk</b>	<b>lwp_newstk(3L)</b>	LWP stack management
<b>lwp_perror</b>	<b>lwp_perror(3L)</b>	LWP error handling
<b>lwp_ping</b>	<b>lwp_status(3L)</b>	LWP status information
<b>lwp_resched</b>	<b>lwp_yield(3L)</b>	control LWP scheduling
<b>lwp_resume</b>	<b>lwp_yield(3L)</b>	control LWP scheduling
<b>lwp_self</b>	<b>lwp_status(3L)</b>	LWP status information
<b>lwp_setpri</b>	<b>lwp_yield(3L)</b>	control LWP scheduling
<b>lwp_setregs</b>	<b>lwp_status(3L)</b>	LWP status information
<b>lwp_setstkcache</b>	<b>lwp_newstk(3L)</b>	LWP stack management
<b>lwp_sleep</b>	<b>lwp_yield(3L)</b>	control LWP scheduling
<b>lwp_stkcswget</b>	<b>lwp_newstk(3L)</b>	LWP stack management
<b>lwp_suspend</b>	<b>lwp_yield(3L)</b>	control LWP scheduling
<b>lwp_yield</b>	<b>lwp_yield(3L)</b>	control LWP scheduling
<b>MINSTACKSZ</b>	<b>lwp_newstk(3L)</b>	LWP stack management
<b>mon_break</b>	<b>mon_create(3L)</b>	LWP routines to manage critical sections
<b>mon_cond_enter</b>	<b>mon_create(3L)</b>	LWP routines to manage critical sections
<b>mon_create</b>	<b>mon_create(3L)</b>	LWP routines to manage critical sections
<b>mon_destroy</b>	<b>mon_create(3L)</b>	LWP routines to manage critical sections
<b>mon_enter</b>	<b>mon_create(3L)</b>	LWP routines to manage critical sections
<b>mon_enumerate</b>	<b>mon_create(3L)</b>	LWP routines to manage critical sections
<b>mon_exit</b>	<b>mon_create(3L)</b>	LWP routines to manage critical sections
<b>mon_waiters</b>	<b>mon_create(3L)</b>	LWP routines to manage critical sections
<b>MONITOR</b>	<b>mon_create(3L)</b>	LWP routines to manage critical sections
<b>msg_enumrecv</b>	<b>msg_send(3L)</b>	LWP send and receive messages
<b>msg_enumsend</b>	<b>msg_send(3L)</b>	LWP send and receive messages
<b>msg_recv</b>	<b>msg_send(3L)</b>	LWP send and receive messages
<b>MSG_RECVALL</b>	<b>msg_send(3L)</b>	LWP send and receive messages
<b>msg_reply</b>	<b>msg_send(3L)</b>	LWP send and receive messages
<b>msg_send</b>	<b>msg_send(3L)</b>	LWP send and receive messages
<b>pod_exit</b>	<b>lwp_create(3L)</b>	LWP thread creation and destruction primitives
<b>pod_getexit</b>	<b>lwp_create(3L)</b>	LWP thread creation and destruction primitives
<b>pod_getmaxpri</b>	<b>pod_getmaxpri(3L)</b>	control LWP scheduling priority
<b>pod_getmaxsize</b>	<b>pod_getmaxpri(3L)</b>	control LWP scheduling priority
<b>pod_setexit</b>	<b>lwp_create(3L)</b>	LWP thread creation and destruction primitives
<b>pod_setmaxpri</b>	<b>pod_getmaxpri(3L)</b>	control LWP scheduling priority
<b>SAMECV</b>	<b>cv_create(3L)</b>	manage LWP condition variables
<b>SAMEMON</b>	<b>mon_create(3L)</b>	LWP routines to manage critical sections
<b>SAMETHREAD</b>	<b>lwp_create(3L)</b>	LWP thread creation and destruction primitives
<b>STKTOP</b>	<b>lwp_newstk(3L)</b>	LWP stack management

## NAME

agt\_create, agt\_enumerate, agt\_trap – map LWP events into messages

## SYNOPSIS

```
#include <lwp/lwp.h>

thread_t agt_create(agt, event, memory)
thread_t *agt;
int event;
caddr_t memory;

int agt_enumerate(vec, maxsize)
thread_t vec[];
int maxsize;

int agt_trap(event)
int event;
```

## DESCRIPTION

Agents are entities that act like threads sending messages when an asynchronous event occurs. **agt\_create()** creates an object called an *agent* which maps the asynchronous event *event* into messages that can be received with **msg\_rcv()** (see **msg\_send(3L)**). *agt* stores the handle on this object. *event* is a UNIX signal number.

**agt\_trap()** causes the event, *event*, to generate an exception (see **exc\_handle(3L)**). Once initialized using **agt\_create()** or **agt\_trap()**, an event can not be remapped to a different style of handling. If traps are enabled, an event will cause the termination of the *thread* running at the time of the trap if the trap exception is not handled. If an exception handler is in place, an exception will be raised. If an agent exists for the event, the event is mapped into a message for the agent. If neither agent nor trap mapping is enabled, the default signal action (SIG\_DFL) is applied to the *pod*. Use of standard UNIX signal handling facilities will defeat the event mapping mechanism.

The message sent by the agent (in the argument buffer) will look like any other message with the sender being the agent. The receive buffer is NULL. A message is always sent by an agent to the thread which created the agent.

All messages sent by an agent contain an **eventinfo\_t**. This structure indicates the thread running at the time the interrupt happened, and the particular event that occurred. Some agent messages contain more information if the particular event warrants it. In this case, a struct containing an **eventinfo\_t** as its first element is passed as the argument buffer. Definitions of these structures are contained in **<lwp/lwp.h>**.

An agent appears to the owning thread just like another thread. It must therefore have some memory for holding its message, as the sender and receiver must belong to the same address space. *memory* is the space an agent will use to store its message. Typically, this is on the stack of the thread that created the agent. It must be of the correct size for the kind of event being created (most events need something to store an **eventinfo\_t**. SIGCHLD events need room for a **sigchld\_t**.)

You should reply to an agent (using **msg\_reply()** (see **msg\_send(3L)**) as you would reply to a thread. Although agents do not ordinarily lose events, the next agent message will not be delivered until a reply is sent to the agent. Thus, an agent appears to the client as an ordinary thread sending messages. An agent will only lose events if the total number of unreplied-to events in a pod exceeds AGENTMEMORY.

**lwp\_destroy()** is used to destroy an agent. All agents created by a thread automatically disappear when that thread dies. **agt\_enumerate()** fills in a list with the ID's of all existing agents and returns the total number of agents. This primitive uses *maxsize* to avoid exceeding the capacity of the list. If the number of agents is greater than *maxsize*, only *maxsize* agents ID's are filled in *vec*. If *maxsize* is zero, **agt\_enumerate()** returns the total number of agents.

The special event `LASTRITES` is caused by the termination of a thread. An agent for `LASTRITES` will be informed about every thread that terminates, regardless of cause. The `eventinfo_code` element of this agent will contain the stack argument that the dead thread was created with. Note: by allocating adjacent space above the thread stack, this argument can be used to point to private information about a thread. The `eventinfo_victimid` element will contain the id of the dead thread.

#### RETURN VALUES

`agt_create()` and `agt_trap()` return:

0        on success.

-1       on failure.

`agt_enumerate()` returns the total number of agents.

#### ERRORS

`agt_trap()` will fail if one or more of the following are true:

`LE_INUSE`            Agent in use for this event.

`LE_INVALIDARG`      Event specified does not exist.

`agt_create()` will fail if one or more of the following are true:

`LE_INUSE`            Trap mapping in use for this event.

`LE_INVALIDARG`      Attempt to create agent for non-existent event.

#### SEE ALSO

`exc_handle(3L)`, `msg_send(3L)`

#### BUGS

Signal handlers always take the `SIG_DFL` action when no agent manages the event.

If a descriptor used by a parent of the pod (such as its standard input) is marked non-blocking by a thread, it should be reset when the pod terminates to prevent the parent from receiving `EWOULDBLOCK` errors on the descriptor. There is no way to prevent this from happening if a pod is terminated with extreme prejudice (for instance, using `SIGKILL`).

If an agent reports that a descriptor has I/O available, there may be more than one occurrence of I/O available from that descriptor. Thus, being informed that `SIGIO` has occurred on socket *s* may mean that there are several messages waiting to be received from *s*. Clients should be careful to clean out all I/O from a descriptor before going back to sleep.

All system calls should be protected with loops testing for `EINTR` (and monitors if multiple threads can try to use system calls concurrently). An `lwp_sleep()` could result in a hidden clock interrupt for example.

#### WARNINGS

`agt_trap()` should not be used for asynchronous events. If an unsuspecting thread which has no exception handler is running at the time of a trapped event, it will be terminated.

Clients should not normally handle signals themselves since the agent mechanism assumes it is the only entity handling signals.

## NAME

`cv_create`, `cv_destroy`, `cv_wait`, `cv_notify`, `cv_broadcast`, `cv_send`, `cv_enumerate`, `cv_waiters`, SAMECV – manage LWP condition variables

## SYNOPSIS

```
#include <lwp/lwp.h>

cv_t cv_create(cv, mid)
cv_t *cv;
mon_t mid;

int cv_destroy(cv)
cv_t cv;

int cv_wait(cv)
cv_t cv;

int cv_notify(cv)
cv_t cv;

int cv_send(cv, tid)
cv_t cv;
lwp_t tid

int cv_broadcast(cv)
cv_t cv;

int cv_enumerate(vec, maxsize)
cv_t vec[ ]; /* will contain list of all conditions */
int maxsize; /* maximum size of vec */

int cv_waiters(cv, vec, maxsize)
cv_t cv; /* condition variable being interrogated */
thread_t vec[ ]; /* which threads are blocked on cv */
int maxsize; /* maximum size of vec */

SAMECV(c1, c2)
```

## DESCRIPTION

Condition variables are useful for synchronization within monitors. By waiting on a condition variable, the currently-held monitor (a condition variable must *always* be used within a monitor) is released atomically and the invoking thread is suspended. When monitors are nested, monitor locks other than the current one are retained by the thread. At some later point, a different thread may awaken the waiting thread by issuing a notification on the condition variable. When the notification occurs, the waiting thread will queue to reacquire the monitor it gave up. It is possible to have different condition variables operating within the same monitor to allow selectivity in waking up threads.

`cv_create()` creates a new condition variable (returned in `cv`) which is bound to the monitor specified by `mid`. It is illegal to access (using `cv_wait()`, `cv_notify()`, `cv_send()` or `cv_broadcast()`) a condition variable from a monitor other than the one it is bound to. `cv_destroy()` removes a condition variable.

`cv_wait()` blocks the current thread and releases the monitor lock associated with the condition (which must also be the monitor lock most recently acquired by the thread). Other monitor locks held by the thread are not affected. The blocked thread is enqueued by its scheduling priority on the condition.

`cv_notify()` awakens at most one thread blocked on the condition variable and causes the awakened thread to queue for access to the monitor released at the time it waited on the condition. It can be dangerous to use `cv_notify()` if there is a possibility that the thread being awakened is one of several threads that are waiting on a condition variable and the awakened thread may not be the one intended. In this case, use of `cv_broadcast()` is recommended.

**cv\_broadcast()** is the same as **cv\_notify()** except that *all* threads blocked on the condition variable are awakened. **cv\_notify()** and **cv\_broadcast()** do nothing if no thread is waiting on the condition. For both **cv\_notify()** and **cv\_broadcast()**, the currently held monitor must agree with the one bound to the condition by **cv\_create()**.

**cv\_send()** is like **cv\_notify()** except that the particular thread **tid** is awakened. If this thread is not currently blocked on the condition, **cv\_send()** reports an error.

**cv\_enumerate()** lists the ID of all of the condition variables. The value returned is the total number of condition variables. The vector supplied is filled in with the ID's of condition variables. **cv\_waiters()** lists the ID's of the threads blocked on the condition variable *cv* and returns the number of threads blocked on *cv*. For both **cv\_enumerate()** and **cv\_waiters()**, *maxsize* is used to avoid exceeding the capacity of the list *vec*. If the number of entries to be filled is greater than *maxsize*, only *maxsize* entries are filled in *vec*. It is legal in both of these primitives to specify a *maxsize* of 0.

SAMECV is a convenient predicate used to compare two condition variables for equality.

#### RETURN VALUES

**cv\_create()**, **cv\_destroy()**, **cv\_send()**, **cv\_wait()**, **cv\_notify()** and **cv\_broadcast()** return:

0           on success.

-1          on failure and set **errno** to indicate the error.

**cv\_enumerate()** returns the total number of condition variables.

**cv\_waiters()** returns the number of threads blocked on a condition variable.

#### ERRORS

**cv\_destroy()** will fail if one or more of the following is true:

LE\_INUSE           Attempt to destroy condition variable being waited on by a thread.

LE\_NONEXIST        Attempt to destroy non-existent condition variable.

**cv\_wait()** will fail if one or more of the following is true:

LE\_NONEXIST        Attempt to wait on non-existent condition variable.

LE\_NOTOWNED        Attempt to wait on a condition without possessing the correct monitor lock.

**cv\_notify()** will fail if one or more of the following is true:

LE\_NONEXIST        Attempt to notify non-existent condition variable.

LE\_NOTOWNED        Attempt to notify condition variable without possessing the correct monitor.

**cv\_send()** will fail if one or more of the following is true:

LE\_NONEXIST        Attempt to awaken non-existent condition variable.

LE\_NOTOWNED        Attempt to awaken condition variable without possessing the correct monitor lock.

LE\_NOWAIT          The specified thread is not currently blocked on the condition.

**cv\_broadcast()** will fail if one or more of the following is true:

LE\_NONEXIST        Attempt to broadcast non-existent condition variable.

LE\_NOTOWNED        Attempt to broadcast condition without possessing the correct monitor lock.

#### SEE ALSO

**mon\_create(3L)**



## NAME

`exc_handle`, `exc_unhandle`, `exc_bound`, `exc_notify`, `exc_raise`, `exc_on_exit`, `exc_uniqpatt` – LWP exception handling

## SYNOPSIS

```
#include <lwp/lwp.h>

int exc_handle(pattern, func, arg)
int pattern;
caddr_t (*func)();
caddr_t arg;

int exc_raise(pattern)
int pattern;

int exc_unhandle()

caddr_t (*exc_bound(pattern, arg))()
int pattern;
caddr_t *arg;

int exc_notify(pattern)
int pattern;

int exc_on_exit(func, arg)
void (*func)();
caddr_t arg;

int exc_uniqpatt()
```

## DESCRIPTION

These primitives can be used to manage exceptional conditions in a thread. Basically, raising an exception is a more general form of non-local goto or *longjmp*, but the invocation is pattern-based. It is also possible to *notify* an exception handler whereby a function supplied by the exception handler is invoked and control is returned to the raiser of the exception. Finally, one can establish a handler which is always invoked upon procedure exit, regardless of whether the procedure exits using a *return* or an exception raised to a handler established prior to the invocation of the exiting procedure.

`exc_handle()` is used to establish an exception handler. `exc_handle()` returns 0 to indicate that a handler has been established. A return of -1 indicates an error in trying to establish the exception handler. If it returns something else, an exception has occurred and any procedure calls deeper than the one containing the handler have disappeared. All exception handlers established by a procedure are automatically discarded when the procedure terminates.

`exc_handle()` binds a *pattern* to the handler, where a pattern is an integer, and two patterns *match* if their values are equal. When an exception is raised with `exc_raise()`, the most recent handler that has established a matching pattern will catch the exception. A special pattern (CATCHALL) is provided which matches any `exc_raise()` pattern. This is useful for handlers which know that there is no chance the resources allocated in a routine can be reclaimed by previous routines in the call chain.

The other two arguments to `exc_handle()` are a function and an argument to that function. `exc_bound()` retrieves these arguments from an `exc_handle()` call made by the specified thread. By using `exc_bound()` to retrieve and call a function bound by the exception handler, a procedure can raise a *notification exception* which allows control to return to the raiser of the exception after the exception is handled.

**exc\_raise()** allows the caller to transfer control (do a non-local goto) to the matching **exc\_handle()**. This matching exception handler is destroyed after the control transfer. At this time, it behaves as if **exc\_handle()** returns with the *pattern* from **exc\_raise()** as the return value. Note: *func* of **exc\_handle()** is not called using **exc\_raise()** — it is only there for notification exceptions. Because the exception handler returns the pattern that invoked it, it is possible for a handler that matches the CATCHALL pattern to *reraise* the exact exception it caught by using **exc\_raise()** on the caught pattern. It is illegal to handle or raise the pattern 0 or the pattern -1. Handlers are searched for pattern matches in the reverse execution order that they are set (i.e., the most recently established handler is searched first).

**exc\_unhandle()** destroys the most recently established exception handler set by the current thread. It is an error to destroy an exit-handler set up by **exc\_on\_exit()**. When a procedure exits, all handlers and exit handlers set in the procedure are automatically deallocated.

**exc\_notify()** is a convenient way to use **exc\_bound**. The function which is bound to *pattern* is retrieved. If the function is not NULL, the function is called with the associated argument and the result is returned. If the function is NULL, **exc\_raise(pattern)** is returned.

**exc\_on\_exit()** specifies an exit procedure and argument to be passed to the exit procedure, which is called when the procedure which sets an exit handler using **exc\_on\_exit()** exits. The exit procedures (more than one may be set) will be called regardless if the setting procedure is exited using a *return* or an **exc\_raise()**. Because the exit procedure is called as if the handling procedure had returned, the argument passed to it should not contain addresses on the handler's stack. However, any value returned by the procedure which established the exit procedure is preserved no matter what the exit procedure returns. This primitive is used in the MONITOR macro to enforce the monitor discipline on procedures.

Some signals can be considered to be synchronous traps. They are usually the starred (\*) signals in the **signal(3V)** man pages. These are: SIGSYS, SIGBUS, SIGEMT, SIGFPE, SIGILL, SIGTRAP, SIGSEGV. If an event is marked as a trap using **agt\_trap()** (see **agt\_create(3L)**) the event will generate exceptions instead of agent messages. This mapping is per-pod, not per-thread. A thread which handles the signal number of one of these as the pattern for **exc\_handle()** will catch such a signal as an exception. The exception will be raised as an **exc\_notify()** so either escape or notification style exceptions can be used, depending on what the matching **exc\_handle()** provides. If the exception is not handled, the thread will terminate. Note: it can be dangerous to supply an exception handler to treat stack overflow since the client's stack is used in raising the exception.

**exc\_uniqpatt()** returns an exception pattern that is not any of the pre-defined patterns (any of the synchronous exceptions or -1 or CATCHALL). Each call to **exc\_uniqpatt()** results in a different pattern. If **exc\_uniqpatt()** cannot guarantee uniqueness, -1 is returned instead the *first* time this happens. Subsequent calls after this error result in patterns which may be duplicates.

#### RETURN VALUES

**exc\_uniqpatt()** returns a unique pattern on success. The *first* time it fails, **exc\_uniqpatt()** returns -1.

**exc\_handle()** returns:

- 0 on success.
- 1 on failure. When **exc\_handle()** returns because of a matching call to **exc\_raise()**, it returns the *pattern* raised by **exc\_raise()**.

On success, **exc\_raise()** transfers control to the matching **exc\_handle()** and does not return. On failure, it returns -1.

**exc\_unhandle()** returns:

- 0 on success.
- 1 on failure.

**exc\_bound()** returns a pointer to a function on success. On failure, it returns NULL.

On success, `exc_notify()` returns the return value of a function, or transfers control to a matching `exc_handle()` and does not return. On failure, it returns `-1`.

`exc_on_exit()` returns `0`.

#### ERRORS

`exc_unhandle()` will fail if one or more of the following is true:

`LE_NONEXIST`            Attempt to remove a non-existent handler.  
                          Attempt to remove an exit handler.

`exc_raise()` will fail if one or more of the following is true:

`LE_INVALIDARG`        Attempt to raise an illegal pattern (`-1` or `0`).  
`LE_NONEXIST`            No context found to raise an exception to.

`exc_handle()` will fail if one or more of the following is true:

`LE_INVALIDARG`        Attempt to handle an illegal pattern (`-1` or `0`).

`exc_uniqpatt()` will fail if one or more of the following is true:

`LE_REUSE`             Possible reuse of existing object. `agt_create(3L)`, `signal(3V)`

#### BUGS

The stack may not contain useful information after an exception has been caught so post-exception debugging can be difficult. The reason for this is that a given handler may call procedures that trash the stack before reraising an exception.

The distinction between traps and interrupts can be problematical.

The environment restored on `exc_raise()` consists of the registers at the time of the `exc_handle()`. As a result, modifications to register variables between the times of `exc_handle()` and `exc_raise()` will not be seen. This problem does not occur in the sun4 implementation.

#### WARNINGS

`exc_on_exit()` passes a simple type as an argument to the exit routine. If you need to pass a complex type, such as `thread_t`, `mon_t`, or `cv_t`, pass a pointer to the object instead.

## NAME

`lwp_create`, `lwp_destroy`, `SAMETHREAD`, `pod_setexit`, `pod_getexit`, `pod_exit` – LWP thread creation and destruction primitives

## SYNOPSIS

```
#include <lwp/lwp.h>
#include <lwp/stackdep.h>

int lwp_create(tid, func, prio, flags, stack, nargs, arg1, ..., argn)
thread_t *tid;
void (*func)();
int prio;
int flags;
stkalign_t *stack;
int nargs;
int arg1, ..., argn;

int lwp_destroy(tid)
thread_t tid;

void pod_setexit(status)
int status;

int pod_getexit(status)
int status;

void pod_exit(status)
int status;

SAMETHREAD(t1, t2)
```

## DESCRIPTION

`lwp_create()` creates a lightweight process which starts at address *func* and has stack segment *stack*. If *stack* is NULL, the thread is created in a suspended state (see below) and no stack or pc is bound to the thread. *prio* is the scheduling priority of the thread (higher priorities are favored by the scheduler). The identity of the new thread is filled in the reference parameter *tid*. *flags* describes some options on the new thread. `LWPSUSPEND` creates the thread in suspended state (see `lwp_yield(3L)`). `LWPNOLASTRITES` will disable the `LASTRITES` agent message when the thread dies. The default (0) is to create the thread in running state with `LASTRITES` reporting enabled. `LWPSEVER` indicates that a thread is only viable as long as non-`LWPSEVER` threads are alive. The pod will terminate if the only living threads are marked `LWPSEVER` and blocked on a lwp resource (for instance, waiting for a message to be sent). *nargs* is the number (0 or more) of simple-type (int) arguments supplied to the thread.

The first time a lwp primitive is used, the lwp library automatically converts the caller (i.e., `main`) into a thread with the highest available scheduling priority (see `pod_getmaxpri(3L)`). The identity of this thread can be retrieved using `lwp_self` (see `lwp_status(3L)`). This thread has the normal SunOS stack given to any *forked* process.

Scheduling is, by default, non-preemptive within a priority, and within a priority, threads enter the run queue on a FIFO basis (that is, whenever a thread becomes eligible to run, it goes to the end of the run queue of its particular priority). Thus, a thread continues to run until it voluntarily relinquishes control or an event (including thread creation) occurs to enable a higher priority thread. Some primitives may cause the current thread to block, in which case the unblocked thread with the highest priority runs next. When several threads are created with the same priority, they are queued for execution in the order of creation. This order may not be preserved as threads yield and block within a priority. If an agent owned by a thread with a higher priority is invoked, that thread will preempt the currently running one.

There is no concept of ancestry in threads: the creator of a thread has no special relation to the thread it created. When all threads have died, the pod terminates.

**lwp\_destroy()** is a way to explicitly terminate a thread or agent (instead of having an executing thread “fall through”, which also terminates the thread). *tid* specifies the id of the thread or agent to be terminated. If *tid* is **SELF**, the invoking thread is destroyed. Upon termination, the resources (messages, monitor locks, agents) owned by the thread are released, in some cases resulting in another thread being notified of the death of its peer (by having a blocking primitive become unblocked with an error indication). A thread may terminate itself explicitly, although self-destruction is automatic when it returns from the procedure specified in the **lwp\_create()** primitive.

**pod\_setexit()** sets the exit status for a pod. This value will be returned to the parent process of the pod when the pod dies (default is 0). **exit(3)** terminates the current *thread*, using the argument supplied to *exit* to set the current value of the exit status. **on\_exit(3)** establishes an action that will be taken when the entire pod terminates. **pod\_exit()** is available to terminate the pod immediately with the final actions established by **on\_exit**. If you wish to terminate the pod immediately, **pod\_exit()** or **exit(2V)** should be used.

**pod\_getexit()** returns the current value of the pod's exit status.

**SAMETHREAD()** is a convenient predicate used to compare two threads for equality.

#### RETURN VALUES

**lwp\_create()**, and **lwp\_destroy()** return:

- 0        on success.
- 1       on failure.

**pod\_getexit()** returns the current exit status of the pod.

#### ERRORS

**lwp\_create()** will fail if one or more of the following are true:

- LE\_ILLPRIO**        Illegal priority.
- LE\_INVALIDARG**    Too many arguments (> 512).
- LE\_NOROOM**        Unable to allocate memory for thread context.

**lwp\_destroy()** will fail if one or more of the following are true:

- LE\_NONEXIST**        Attempt to destroy a thread or agent that does not exist.

#### SEE ALSO

**exit(2V)**, **exit(3)**, **lwp\_yield(3L)**, **on\_exit(3)**, **pod\_getmaxpri(3L)**

#### WARNINGS

Some special threads may be created silently by the lwp library. These include an *idle* thread that runs when no other activity is going on, and a *reaper* thread that frees stacks allocated by **lwp\_newstk**. These special threads will show up in status calls. A pod will terminate if these special threads are the only ones extant.

## NAME

`lwp_ctxinit`, `lwp_ctxremove`, `lwp_ctxset`, `lwp_ctxmemget`, `lwp_ctxmemset`, `lwp_fpset`, `lwp_libcset` – special LWP context operations

## SYNOPSIS

```
#include <lwp/lwp.h>

int lwp_ctxset(save, restore, ctxsize, optimize)
void (*save)(/* caddr_t ctx, thread_t old, thread_t new */);
void (*restore)(/* caddr_t ctx, thread_t old, thread_t new */);
unsigned int ctxsize;
int optimize;

int lwp_ctxinit(tid, cookie)
thread_t tid;          /* thread with special contexts */
int cookie;           /* type of context */

int lwp_ctxremove(tid, cookie)
thread_t tid;
int cookie;

int lwp_ctxmemget(mem, tid, ctx)
caddr_t mem;
thread_t tid;
int ctx;

int lwp_ctxmemset(mem, tid, ctx)
caddr_t mem;
thread_t tid;
int ctx;

int lwp_fpset(tid)
thread_t tid;          /* thread utilizing floating point hardware */

int lwp_libcset(tid)
thread_t tid;          /* thread utilizing errno */
```

## DESCRIPTION

Normally on a context switch, only machine registers are saved/restored to provide each thread its own virtual machine. However, there are other hardware and software resources which can be multiplexed in this way. For example, floating point registers can be used by several threads in a pod. As another example, the global value `errno` in the standard C library may be used by all threads making system calls.

To accommodate the variety of contexts that a thread may need without requiring all threads to pay for unneeded switching overhead, `lwp_ctxinit()` is provided. This primitive allows a client to specify that a given thread requires certain context to be saved and restored across context switches (by default just the machine registers are switched). More than one special context may be given to a thread.

To use `lwp_ctxinit()`, it is first necessary to define a special context. `lwp_ctxset()` specifies save and restore routines, as well as the size of the context that will be used to hold the switchable state. The *save* routine will automatically be invoked when an active thread is blocked and the *restore* routine will be invoked when a blocked thread is restarted. These routines will be passed a pointer to a buffer (initialized to all 0's) of size *ctxsize* which is allocated by the LWP library and used to hold the volatile state. In addition, the identity of the thread whose special context is being saved (old) and the identity of the thread being restarted (new) are passed in to the *save* and *restore* routines. `lwp_ctxset()` returns a cookie used by subsequent `lwp_ctxinit()` calls to refer to the kind of context just defined. If the *optimize* flag is TRUE, a special context switch action will not be invoked unless the thread resuming execution differs from the last thread to use the special context and also uses the special context. If the *optimize* flag is FALSE, the *save* routine will always be invoked immediately when the thread using this context is scheduled out and the *restore* routine will be invoked immediately when a new thread using this context is scheduled in. Note

that an unoptimized special context is protected from threads which do not use the special context but which do affect the context state. `lwp_ctxremove()` can be used to remove a special context installed by `lwp_ctxinit()`.

Because context switching is done by the scheduler on behalf of a thread, it is an error to use an LWP primitive in an action done at context switch time. Also, the stack used by the save and restore routines belongs to the scheduler, so care should be taken not to use lots of stack space. As a result of these restrictions, only knowledgeable users should write their own special context switching routines.

`lwp_ctxmemget()` and `lwp_ctxmemset()` are used to retrieve and set (respectively) the memory associated with a given special context (*ctx*) and a given thread (*tid*). *mem* is the address of client memory that will hold the context information being retrieved or set. Note that the special context *save* and *restore* routines may be NULL, so pure data may be associated with a given thread using these primitives.

Several kinds of special contexts are predefined. To allow a thread to share floating point hardware with other threads, the `lwp_fpset()` primitive is available. The floating-point hardware bound at compile-time is selected automatically. To multiplex the global variable `errno`, `lwp_libcset()` is used to have `errno` become part of the context of thread *tid*.

Special contexts can be used to assist in managing stacks. See `lwp_newstk(3L)` for details.

#### RETURN VALUES

On success, `lwp_ctxset()` returns a cookie to be used by subsequent calls to `lwp_ctxinit()`. If unable to define the context, it returns `-1`.

#### ERRORS

`lwp_ctxinit()` will fail if one or more of the following are true:

`LE_INUSE`                    This special context already set for this thread.

`lwp_ctxremove()` will fail if one or more of the following are true:

`LE_NONEXIST`                The specified context is not set for this thread.

`lwp_ctxset()` will fail if one or more of the following are true:

`LE_NOROOM`                  Unable to allocate memory to define special context.

#### SEE ALSO

`lwp_newstk(3L)`

#### BUGS

The floating point contexts should be initialized implicitly for those threads that use floating point.

**NAME**

`lwp_checkstkset`, `lwp_stkcswset`, `CHECK`, `lwp_setstkcache`, `lwp_newstk`, `lwp_datastk`, `STKTOP` – LWP stack management

**SYNOPSIS**

```
#include <lwp/lwp.h>
#include <lwp/check.h>
#include <lwp/lwpmachdep.h>
#include <lwp/stackdep.h>

CHECK(location, result)

int lwp_checkstkset(tid, limit)
thread_t tid;
caddr_t limit;

int lwp_stkcswset(tid, limit)
thread_t tid;
caddr_t limit;

int lwp_setstkcache(minstksz, numstks)
int minstksz;
int numstks;

stkalign_t *lwp_newstk()

stkalign_t *lwp_datastk(data, size, addr)
caddr_t data;
int size;
caddr_t *addr;

STKTOP(s)
```

**DESCRIPTION**

Stacks are problematical with lightweight processes. What is desired is that stacks for each thread are redzone protected so that one thread's stack does not unexpectedly grow into the stack of another. In addition, stacks should be of infinite length, grown as needed. The process stack is a maximum-sized segment (see `getrlimit(2)`.) This stack is redzone protected, and you can even try to extend it beyond its initial maximum size in some cases. With SunOS 4.x, it is possible to efficiently allocate large stacks that have red zone protection, and the LWP library provides some support for this. For those systems that do not have flexible memory management, the LWP library provides assistance in dealing with the problems of maintaining multiple stacks.

The stack used by `main()` is the same stack that the system allocates for a process on `fork(2V)`. For allocating other thread stacks, the client is free to use any statically or dynamically allocated memory (using memory from `main()`'s stack is subject to the stack resource limit for any process created by `fork()`). In addition, the `LASTRITES` agent message is available to free allocated resources when a thread dies. The size of any stack should be at least `MINSTACKSZ * sizeof(stkalign_t)`, because the LWP library will use the client stack to execute primitives. For very fast dynamically allocated stacks, a stack cacheing mechanism is available. `lwp_setstkcache()` allocates a cache of stacks. Each time the cache is empty, it is filled with `numstks` new stacks, each containing at least `minstksz` bytes. `minstksz` will automatically be augmented to take into account the stack needs of the LWP library. `lwp_newstk()` returns a cached stack that is suitable for use in an `lwp_create()` call. `lwp_setstkcache()` must be called (once) prior to any use of `lwp_newstk`. If running under SunOS 4.x, the stacks allocated by `lwp_newstk()` will be red-zone protected (an attempt to reference below the stack bottom will result in a `SIGSEGV` event).

Threads created with stacks from `lwp_newstk()` should not use the `NOLASTRITES` flag. If they do, cached stacks will not be returned to the cache when a thread dies.



**lwp\_datastk()** also returns a red-zone protected stack like **lwp\_newstk()** does. It copies any amount of data (subject to the size limitations imposed by **lwp\_setstkcache**) onto the stack *above* the stack top that it returns. *data* points to information of *size* bytes to be copied. The exact location where the data is stored is returned in the reference parameter *addr*. Because **lwp\_create()** only passes simple types to the newly-created thread, **lwp\_datastk()** is useful to pass a more complex argument: Call **lwp\_datastk()** to get an initialized stack, and pass the address of the data structure (*addr*) as an argument to the new thread.

A *reaper* thread running at the maximum pod priority is created by **lwp\_setstkcache**. It's action may be delayed by other threads running at that priority, so it is suggested that the maximum pod priority not be used for client-created threads when **lwp\_newstk()** is being used. Altering the maximum pod priority with **pod\_setmaxpri()** will have the side effect of increasing the reaper thread priority as well.

The stack address passed to **lwp\_create()** represents the top of the stack: the LWP library will not use any addresses at or above it. Thus, it is safe to store information above the stack top if there is room there.

For stacks that are not protected with hardware redzones, some protection is still possible. For any thread *tid* with stack boundary *limit* made part of a special context with **lwp\_checkstkset()**, the CHECK macro may be used. This macro, if used at the beginning of each procedure (and before local storage is initialized (it is all right to *declare* locals though)), will check that the stack limit has not been violated. If it has, the non-local *location* will be set to *result* and the procedure will return. CHECK is not perfect, as it is possible to call a procedure with many arguments after CHECK validates the stack, only to have these arguments clobber the stack before the new procedure is entered.

**lwp\_stkcswset()** checks at context-switch time the stack belonging to thread *tid* for passing stack boundary *limit*. In addition, a checksum at the bottom of the stack is validated to ensure that the stack did not temporarily grow beyond its limit. This is automated and more efficient than using CHECK, but by the time a context switch occurs, it's too late to do much but **abort(3)** if the stack was clobbered.

To portably use statically allocated stacks, the macros in `<lwp/stackdep.h>` should be used. Declare a stack *s* to be an array of `stkaligned_t`, and pass the stack to **lwp\_create()** as STKTOP(*s*).

#### RETURN VALUES

**lwp\_checkstkset()** and **lwp\_stkcswset()** return 0.

**lwp\_setstkcache()** returns the actual size of the stacks allocated in the cache.

**lwp\_newstk()** and **lwp\_datastk()** return a valid new stack address on success. On failure, they return 0.

#### SEE ALSO

**getrlimit(2)**, **abort(3)**

#### WARNINGS

**lwp\_datastk()** should not be directly used in a **lwp\_create()** call since C does not guarantee the order in which arguments to a function are evaluated.

#### BUGS

C should provide support for heap-allocated stacks at procedure entry time. The hardware should be segment-based to eliminate the problem altogether.

**NAME**

`lwp_geterr`, `lwp_perror`, `lwp_errstr` – LWP error handling

**SYNOPSIS**

```
#include <lwp/lwp.h>
#include <lwp/lwperror.h>

lwp_err_t lwp_geterr();

void
lwp_perror(s)
char *s;

char **lwp_errstr();
```

**DESCRIPTION**

When a primitive fails (returns `-1`), `lwp_geterr()` can be used to obtain the identity of the error (which is part of the context for each lwp). `lwp_perror()` can be used to print an error message on the standard error file (analogous to `perror(3)`) when a lwp primitive returns an error indication. `lwp_perror()` uses the same mechanism as `lwp_geterr()` to obtain the last error. `lwp_errstr` returns a pointer to the (NULL-terminated) list of error messages.

`lwp_libcset` (see `lwp_ctxinit(3L)`) allows `errno` from the standard C library reflect a per-thread value rather than a per-pod value.

**SEE ALSO**

`lwp_ctxinit(3L)`, `perror(3)`

**NAME**

`lwp_self`, `lwp_ping`, `lwp_enumerate`, `lwp_getstate`, `lwp_setregs`, `lwp_getregs` – LWP status information

**SYNOPSIS**

```
#include <lwp/lwp.h>
#include <lwp/lwpmachdep.h>

int
lwp_enumerate(vec, maxsize)
thread_t vec[]; /* list of id's to be filled in */
int maxsize;    /* number of elements in vec */

int
lwp_ping(tid)
thread_t tid;

int
lwp_getregs(tid, machstate)
thread_t tid;
machstate_t *machstate;

int
lwp_setregs(tid, machstate)
thread_t tid;
machstate_t *machstate;

int
lwp_getstate(tid, statvec)
thread_t tid;
statvec_t *statvec;

int
lwp_self(tid)
thread_t *tid;
```

**DESCRIPTION**

`lwp_self()` returns the ID of the current thread in *tid*. This is the *only* way to retrieve the identity of *main*.

`lwp_enumerate()` fills in a list with the ID's of all existing threads and returns the total number of threads. This primitive will use *maxsize* to avoid exceeding the capacity of the list. If the number of threads is greater than *maxsize*, only *maxsize* thread ID's are filled in *vec*. If *maxsize* is zero, `lwp_enumerate()` just returns the total number of threads.

`lwp_getstate()` is used to retrieve the context of a given thread. It is possible to see what object (thread, monitor, etc.) if any that thread is blocked on, and the scheduling priority of the thread.

`lwp_ping` returns 0 (no error) if the thread *tid* exists. Otherwise, -1 is returned.

`lwp_setregs` sets the machine-dependent context (i.e., registers) of a thread. The next time the thread is scheduled in, this context is installed. Consult `lwpmachdep.h` for the details. `lwp_getregs` retrieves the machine-dependent context. Note: the registers may not be meaningful unless the thread in question is blocked or suspended because the state of the registers as of the most recent context switch is returned.

**RETURNS**

Upon successful completion, `lwp_self` and `lwp_getstate()` return 0, -1 on error.

`lwp_enumerate()` returns the total number of threads.

`lwp_ping` returns 0 if the specified thread exists, else -1.

**ERRORS**

`lwp_getstatea()`, `lwp_ping()`, and `lwp_setstate()` will fail if one or more of the following is true:

`LE_NONEXIST`            Attempt to get the status of a non-existent thread.

## NAME

`lwp_yield`, `lwp_suspend`, `lwp_resume`, `lwp_join`, `lwp_setpri`, `lwp_resched`, `lwp_sleep` – control LWP scheduling

## SYNOPSIS

```
#include <lwp/lwp.h>

int lwp_yield(tid)
thread_t tid;

int lwp_sleep(timeout)
struct timeval *timeout;

int lwp_resched(prio)
int prio;

int lwp_setpri(tid, prio)
thread_t tid;
int prio;

int lwp_suspend(tid)
thread_t tid;

int lwp_resume(tid)
thread_t tid;

int lwp_join(tid)
thread_t tid;
```

## DESCRIPTION

`lwp_yield()` allows the currently running thread to voluntarily relinquish control to another thread *with the same scheduling priority*. If `tid` is `SELF`, the next thread in the same priority queue of the yielding thread will run and the current thread will go the end of the scheduling queue. Otherwise, it is the ID of the thread to run next, and the current thread will take second place in the scheduling queue.

`lwp_sleep()` blocks the thread executing this primitive for at least the time specified by `timeout`.

Scheduling of threads is, by default, preemptive (higher priorities preempt lower ones) across priorities and non-preemptive within a priority. `lwp_resched()` moves the front thread for a given priority to the end of the scheduling queue. Thus, to achieve a preemptive round-robin scheduling discipline, a high priority thread can periodically wake up and shuffle the queue of threads at a lower priority. `lwp_resched()` does not affect threads which are blocked. If the priority of the rescheduled thread is the same as that of the caller, the effect is the same as `lwp_yield()`.

`lwp_setpri()` is used to alter (raise or lower) the scheduling priority of the specified thread. If `tid` is `SELF`, the priority of the invoking thread is set. Note: if the priority of the affected thread becomes greater than that of the caller and the affected thread is not blocked, the caller will not run next. `lwp_setpri()` can be used on either blocked or unblocked threads.

`lwp_join()` blocks the thread issuing the join until the thread `tid` terminates. More than one thread may join `tid`.

`lwp_suspend()` makes the specified thread ineligible to run. If `tid` is `SELF`, the caller is itself suspended. `lwp_resume()` undoes the effect of `lwp_suspend()`. If a blocked thread is suspended, it will not run until it has been unblocked as well as explicitly made eligible to run using `lwp_resume()`. By suspending a thread, one can safely examine it without worrying that its execution-time state will change.

## NOTES

When scheduling preemptively, be sure to use monitors to protect shared data structures such as those used by the standard I/O library.

**RETURN VALUES**

**lwp\_yield()**, **lwp\_sleep()**, **lwp\_resched()**, **lwp\_join()**, **lwp\_suspend()** and **lwp\_resume()** return:

0 on success.

-1 on failure.

**lwp\_setpri()** returns the previous priority on success. On failure, it returns -1.

**ERRORS**

**lwp\_yield()** will fail if one or more of the following is true:

**LE\_ILLPRIO** Attempt to yield to thread with different priority.

**LE\_INVALIDARG** Attempt to yield to a blocked thread.

**LE\_NONEXIST** Attempt to yield to a non-existent thread.

**lwp\_sleep()** will fail if one or more of the following is true:

**LE\_INVALIDARG** Illegal timeout specified.

**lwp\_resched()** will fail if one or more of the following is true:

**LE\_ILLPRIO** The priority queue specified contains no threads to reschedule.

**LE\_INVALIDARG** Attempt to reschedule thread at priority greater than that of the caller.

**lwp\_setpri()** will fail if one or more of the following is true:

**LE\_INVALIDARG** The priority specified is beyond the maximum available to the pod.

**LE\_NONEXIST** Attempt to set priority of a non-existent thread.

**lwp\_join()** will fail if one or more of the following are true:

**LE\_NONEXIST** Attempt to join a thread that does not exist.

**lwp\_suspend()** will fail if one or more of the following is true:

**LE\_NONEXIST** Attempt to suspend a non-existent thread.

**lwp\_resume()** will fail if one or more of the following is true:

**LE\_NONEXIST** Attempt to resume a non-existent thread.

## NAME

`mon_create`, `mon_destroy`, `mon_enter`, `mon_exit`, `mon_enumerate`, `mon_waiters`, `mon_cond_enter`, `mon_break`, `MONITOR`, `SAMEMON` – LWP routines to manage critical sections

## SYNOPSIS

```
#include <lwp/lwp.h>

int mon_create(mid)
mon_t *mid;

int mon_destroy(mid)
mon_t mid;

int mon_enter(mid)
mon_t mid;

int mon_exit(mid)
mon_t mid;

int mon_enumerate(vec, maxsize)
mon_t vec[]; /* list of all monitors */
int maxsize; /* max size of vec */

int mon_waiters(mid, owner, vec, maxsize)
mon_t mid; /* monitor in question */
thread_t *owner; /* which thread owns the monitor */
thread_t vec[]; /* list of blocked threads */
int maxsize; /* max size of vec */

int mon_cond_enter(mid)
mon_t mid;

int mon_break(mid)
mon_t mid;

void MONITOR(mid)
mon_t mid;

int SAMEMON(m1, m2)
mon_t m1;
mon_t m2;
```

## DESCRIPTION

Monitors are used to synchronize access to common resources. Although it is possible (on a uniprocessor) to use knowledge of how scheduling priorities work to serialize access to a resource, monitors (and condition variables) provide a general tool to provide the necessary synchronization.

`mon_create()` creates a new monitor and returns its identity in *mid*. `mon_destroy()` destroys a monitor, as well as any conditions bound to it (see `cv_create(3L)`). Because the lifetime of a monitor can transcend the lifetime of the LWP that created it, monitor destruction is not automatic upon LWP destruction.

`mon_enter()` blocks the calling thread (if the monitor is in use) until the monitor becomes free by being exited or by waiting on a condition (see `cv_create(3L)`). Threads unable to gain entry into the monitor are queued for monitor service by the priority of the thread requesting monitor access, FCFS within a priority. Monitor calls may nest. If, while holding monitor M1 a request for monitor M2 is made, M1 will be held until M2 can be acquired.

`mon_cond_enter()` will enter the monitor only if the monitor is not busy. Otherwise, an error is returned.

`mon_enter()` and `mon_cond_enter()` will allow a thread which already has the monitor to reenter the monitor. In this case, the nesting level of monitor entries is returned. Thus, the first time a monitor is entered, `mon_enter()` returns 0. The next time the monitor is entered, `mon_enter()` returns 1. `mon_exit()` frees the current monitor and allows the next thread blocked on the monitor (if any) to enter

the monitor. However, if a monitor is entered more than once, **mon\_exit()** returns the previous monitor nesting level without freeing the monitor to other threads. Thus, if the monitor was not reentered, **mon\_exit()** returns 0.

**mon\_enumerate()** lists all the monitors in the system. The vector supplied is filled in with the ID's of the monitors. *maxsize* is used to avoid exceeding the capacity of the list. If the number of monitors is greater than *maxsize*, only *maxsize* monitor ID's are filled in *vec*.

**mon\_waiters()** puts the thread that currently owns the monitor in *owner* and all threads blocked on the monitor in *vec* (subject to the *maxsize* limitation), and returns the number of waiting threads.

**mon\_break()** forces the release of a monitor lock not necessarily held by the invoking thread. This enables the next thread blocked on the monitor to enter it.

**MONITOR** is a macro that can be used at the start of a procedure to indicate that the procedure is a monitor. It uses the exception handling mechanism to ensure that the monitor is exited automatically when the procedure exits. Ordinarily, this single macro replaces paired **mon\_enter()**-**mon\_exit()** calls in a monitor procedure.

The **SAMEMON** macro is a convenient predicate used to compare two monitors for equality.

Monitor locks are released automatically when the LWP holding them dies. This may have implications for the validity of the monitor invariant (a condition that is always true *outside* of the monitor) if a thread unexpectedly terminates.

#### RETURN VALUES

**mon\_create()** returns the ID of a new monitor.

**mon\_destroy()** returns:

0        on success.

-1       on failure.

**mon\_enter()** returns the nesting level of the monitor.

**mon\_exit()** returns the previous nesting level on success. On failure, it returns -1.

**mon\_enumerate()** returns the total number of monitors.

**mon\_waiters()** returns the number of threads waiting for the monitor.

**mon\_cond\_enter()** returns the nesting level of the monitor if the monitor is not busy. If the monitor is busy, it returns -1.

**mon\_break()** returns:

0        on success.

-1       on failure.

The macro **SAMEMON()** returns 1 if the monitors specified by *m1* and *m2* are equal. It returns 0 otherwise.

#### ERRORS

**mon\_break()** will fail if one or more of the following are true:

LE\_NONEXIST        Attempt to break lock on non-existent monitor.

LE\_NOTOWNED        Attempt to break a monitor lock that is not set.

**mon\_cond\_enter()** will fail if one or more of the following are true:

LE\_INUSE            The requested monitor is being used by another thread.

LE\_NONEXIST        Attempt to destroy non-existent monitor.

**mon\_destroy()** will fail if one or more of the following are true:

LE\_INUSE            Attempt to destroy a monitor that has threads blocked on it.

LE\_NONEXIST        Attempt to destroy non-existent monitor.

**mon\_exit()** will fail if one or more of the following are true:

LE\_INVALIDARG     Attempt to exit a monitor that the thread does not own.

LE\_NONEXIST        Attempt to exit non-existent monitor.

**SEE ALSO**

**cv\_create(3L)**

**BUGS**

There should be language support to enforce the monitor enter-exit discipline.



## NAME

`msg_send`, `msg_rcv`, `msg_reply`, `MSG_RECVALL`, `msg_enumsend`, `msg_enumrcv` – LWP send and receive messages

## SYNOPSIS

```
#include <lwp/lwp.h>

int msg_send(dest, arg, argsize, res, ressize)
thread_t dest; /* destination thread */
caddr_t arg; /* argument buffer */
int argsize; /* size of argument buffer */
caddr_t res; /* result buffer */
int ressize; /* size of result buffer */

int msg_rcv(sender, arg, argsize, res, ressize, timeout)
thread_t *sender; /* value-result: sending thread or agent */
caddr_t *arg; /* argument buffer */
int *argsize; /* argument size */
caddr_t *res; /* result buffer */
int *ressize; /* result size */
struct timeval *timeout; /* POLL, INFINITY, else timeout */

int msg_reply(sender)
thread_t sender; /* agent id or thread id */

int msg_enumsend(vec, maxsize)
thread_t vec[]; /* list of blocked senders */
int maxsize;

int msg_enumrcv(vec, maxsize)
thread_t vec[]; /* list of blocked receivers */
int maxsize;

int MSG_RECVALL(sender, arg, argsize, res, ressize, timeout)
thread_t *sender;
caddr_t *arg;
int *argsize;
caddr_t *res;
int *ressize;
struct timeval *timeout;
```

## DESCRIPTION

Each thread queues messages addressed to it as they arrive. Threads may either specify that a particular sender's message is to be received next, or that *any* sender's message may be received next.

`msg_send()` specifies a message buffer and a reply buffer, and initiates one half of a rendezvous with the receiver. The sender will block until the receiver replies using `msg_reply()`. `msg_rcv()` initiates the other half of a rendezvous and blocks the invoking thread until a corresponding `msg_send()` is received. When unblocked by `msg_send()`, the receiver may read the message and generate a reply by filling in the reply buffer and issuing `msg_reply()`. `msg_reply()` unblocks the sender. Once a reply is sent, the receiver should no longer access either the message or reply buffer.

In `msg_send()`, `argsize` specifies the size in bytes of the argument buffer `argbuf`, which is intended to be a read-only (to the receiver) buffer. `ressize` specifies the size in bytes of the result buffer `resbuf`, which is intended to be a write-only (to the receiver) buffer. `dest` is the thread that is the target of the send.

`msg_rcv()` blocks the receiver until:

- A message from the agent or thread bound to *sender* has been sent to the receiver or,
- *sender* points to a `THREADNULL`-valued variable and *any* message has been sent to the receiver from a thread or agent, or,
- After the time specified by *timeout* elapses and no message is received.

If *timeout* is `POLL`, `msg_rcv()` returns immediately, returning success if the message expected has arrived; otherwise an error is returned. If *timeout* is `INFINITY`, `msg_rcv()` blocks forever or until the expected message arrives. If *timeout* is any other value `msg_rcv()` blocks for the time specified by *timeout* or until the expected message arrives, whichever comes first. When `msg_rcv()` returns, *sender* is filled in with the identity of the sending thread or agent, and the buffer addresses and sizes specified by the matching send are stored in *arg*, *argsize*, *res*, and *ressize*.

`msg_enumsend()` and `msg_enumrcv()` are used to list all of the threads blocked on sends (awaiting a reply) and receives (awaiting a send), respectively. The value returned is the number of such blocked threads. The vector supplied by the client is filled in (subject to the *maxsize* limitation) with the ID's of the blocked threads. *maxsize* is used to avoid exceeding the capacity of the list. If the number of threads blocked on sends or receives is greater than *maxsize*, only *maxsize* thread ID's are filled in *vec*. If *maxsize* is 0, just the total number of blocked threads is returned.

*sender* in `msg_rcv()` is a reference parameter. If you wish to receive from *any* sender, be sure to reinitialize the thread *sender* points to as `THREADNULL` before each use (do not use the address of `THREADNULL` for the sender). Alternatively, use the `MSG_RECVALL()` macro. This macro has the same parameters as `msg_rcv()`, but ensures that the sender is properly initialized to allow receipt from any sender. `MSG_RECVALL()` returns the result from `msg_rcv`.

#### RETURN VALUES

`msg_send()`, `msg_rcv()`, `MSG_RECVALL()` and `msg_reply()` return:

- 0 on success.
- 1 on failure.

`msg_enumsend()` returns the number of threads blocked on `msg_send()`.

`msg_enumrcv()` returns the number of threads blocked on `msg_rcv()`.

#### ERRORS

`msg_rcv()` will fail if one or more of the following is true:

`LE_INVALIDARG` An illegal timeout was specified.  
The sender address is that of `THREADNULL`.

`LE_NONEXIST` The specified thread or agent does not exist.

`LE_TIMEOUT` Timed out before message arrived.

`msg_reply()` will fail if one or more of the following is true:

`LE_NONEXIST` Attempt to reply to a sender that does not exist or has terminated.

`LE_NOWAIT` Attempt to reply to a sender that is not expecting a reply.

`msg_send()` will fail if one or more of the following is true:

`LE_INVALIDARG` Attempt to send a message to yourself.

`LE_NONEXIST` The specified destination thread does not exist or has terminated.

**NAME**

`pod_getmaxpri`, `pod_getmaxsize`, `pod_setmaxpri` – control LWP scheduling priority

**SYNOPSIS**

```
int pod_getmaxpri()
int pod_getmaxsize()
int pod_setmaxpri(maxprio)
int maxprio;
```

**DESCRIPTION**

The LWP library is self-initializing: the first time you use a primitive that requires threads to be supported, *main* is automatically converted into a thread. A pod will terminate when all client-created lightweight threads (including the thread bound to *main*) are dead.

By default, only a single priority (MINPRIO) is available. However, by using `pod_setmaxpri()`, you can make an arbitrary number (up to the limit imposed by the implementation) of priorities available. The *main* thread will receive the highest available scheduling priority at the time of initialization. By using `pod_setmaxpri()` before any other LWP primitives, you can ensure that *main* will receive the same priority as the argument to `pod_setmaxpri()`. `pod_setmaxpri()` can be called repeatedly, as long as the number of scheduling priorities (*maxprio*) increases with each call.

`pod_getmaxpri()` returns the current number of available priorities. Priorities are numbered from 1 (MINPRIO) to MAXPRIO.

The implementation-dependent maximum number of priorities available can be retrieved using `pod_getmaxsize()`. This value will never be less than 255.

**RETURN VALUES**

`pod_getmaxpri()` returns the number of priority levels set by the most recent call to `pod_setmaxpri()`.

`pod_getmaxsize()` returns the maximum number of priorities your system supports.

`pod_setmaxpri()` returns:

0        on success.  
-1       on failure.

**ERRORS**

`pod_setmaxpri()` will fail if one or more of the following are true:

LE_INVALIDARG	Attempt to allocate more priorities than supported.
LE_NOROOM	No internal memory left to create pod.



**NAME**

intro – introduction to mathematical library functions and constants

**SYNOPSIS**

```
#include <sys/ieeefp.h>
#include <floatingpoint.h>
#include <math.h>
```

**DESCRIPTION**

The include file `<math.h>` contains declarations of all the functions described in Section 3M that are implemented in the math library, `libm`. C programs should be linked with the `-lm` option in order to use this library.

`<sys/ieeefp.h>` and `<floatingpoint.h>` define certain types and constants used for `libm` exception handling, conforming to ANSI/IEEE Std 754-1985, the *IEEE Standard for Binary Floating-Point Arithmetic*.

**ACKNOWLEDGEMENT**

The Sun version of `libm` is based upon and developed from ideas embodied and codes contained in 4.3 BSD, which may not be compatible with earlier BSD or UNIX implementations.

**IEEE ENVIRONMENT**

The IEEE Standard specifies modes for rounding direction, precision, and exception trapping, and status reflecting accrued exceptions. These modes and status constitute the IEEE run-time environment. On Sun-2 and Sun-3 systems without 68881 floating-point co-processors, only the default rounding direction to nearest is available, only the default non-stop exception handling is available, and accrued exception bits are not maintained.

**IEEE EXCEPTION HANDLING**

The IEEE Standard specifies exception handling for `aint`, `ceil`, `floor`, `irint`, `remainder`, `rint`, and `sqrt`, and suggests appropriate exception handling for `fp_class`, `copysign`, `fabs`, `finite`, `fmod`, `isinf`, `isnan`, `ilogb`, `ldexp`, `logb`, `nextafter`, `scalb`, `scalbn` and `signbit`, but does not specify exception handling for the other `libm` functions.

For these other unspecified functions the spirit of the IEEE Standard is generally followed in `libm` by handling invalid operand, singularity (division by zero), overflow, and underflow exceptions, as much as possible, in the same way they are handled for the fundamental floating-point operations such as addition and multiplication.

These unspecified functions are usually not quite correctly rounded, may not observe the optional rounding directions, and may not set the inexact exception correctly.

**SYSTEM V EXCEPTION HANDLING**

The *System V Interface Definition* (SVID) specifies exception handling for some `libm` functions: `j0()`, `j1()`, `jn()`, `y0()`, `y1()`, `yn()`, `exp()`, `log()`, `log10()`, `pow()`, `sqrt()`, `hypot()`, `lgamma()`, `sinh()`, `cosh()`, `sin()`, `cos()`, `tan()`, `asin()`, `acos()`, and `atan2()`. See `matherr(3M)` for a discussion of the extent to which Sun's implementation of `libm` follows the SVID when it is consistent with the IEEE Standard and with hardware efficiency.

**LIST OF MATH LIBRARY FUNCTIONS**

Name	Appears on Page	Description
–	<code>bessel(3M)</code>	Bessel functions
–	<code>frexp(3M)</code>	floating-point analysis
–	<code>hyperbolic(3M)</code>	hyperbolic functions
–	<code>ieee_functions(3M)</code>	IEEE classification
–	<code>ieee_test(3M)</code>	IEEE tests for compliance
–	<code>ieee_values(3M)</code>	returns double-precision IEEE infinity
–	<code>trig(3M)</code>	trigonometric functions
<code>acos</code>	<code>trig(3M)</code>	trigonometric functions

acosh	hyperbolic(3M)	hyperbolic functions
aint	rint(3M)	round to integral value in floating-point or integer format
anint	rint(3M)	round to integral value in floating-point or integer format
annuity	exp(3M)	exponential, logarithm, power
asin	trig(3M)	trigonometric functions
asinh	hyperbolic(3M)	hyperbolic functions
atan	trig(3M)	trigonometric functions
atan2	trig(3M)	trigonometric functions
atanh	hyperbolic(3M)	hyperbolic functions
cbrt	sqrt(3M)	cube root, square root
ceil	rint(3M)	round to integral value in floating-point or integer format
compound	exp(3M)	exponential, logarithm, power
copysign	ieee_functions(3M)	miscellaneous functions for IEEE arithmetic
cos	trig(3M)	trigonometric functions
cosh	hyperbolic(3M)	hyperbolic functions
erf	erf(3M)	error functions
erfc	erf(3M)	error functions
exp	exp(3M)	exponential, logarithm, power
exp2	exp(3M)	exponential, logarithm, power
exp10	exp(3M)	exponential, logarithm, power
expm1	exp(3M)	exponential, logarithm, power
fabs	ieee_functions(3M)	miscellaneous functions for IEEE arithmetic
finite	ieee_functions(3M)	miscellaneous functions for IEEE arithmetic
floor	rint(3M)	round to integral value in floating-point or integer format
fmod	ieee_functions(3M)	miscellaneous functions for IEEE arithmetic
fp_class	ieee_functions(3M)	miscellaneous functions for IEEE arithmetic
frexp	frexp(3M)	traditional UNIX functions
HUGE	ieee_values(3M)	functions that return extreme values of IEEE arithmetic
HUGE_VAL	ieee_values(3M)	functions that return extreme values of IEEE arithmetic
hypot	hypot(3M)	Euclidean distance
ieee_flags	ieee_flags(3M)	mode and status function for IEEE standard arithmetic
ieee_functions	ieee_functions(3M)	miscellaneous functions for IEEE arithmetic
ieee_handler	ieee_handler(3M)	IEEE exception trap handler function
ieee_test	ieee_test(3M)	IEEE test functions for verifying standard compliance
ieee_values	ieee_values(3M)	functions that return extreme values of IEEE arithmetic
ilogb	ieee_functions(3M)	miscellaneous functions for IEEE arithmetic
infinity	ieee_values(3M)	functions that return extreme values of IEEE arithmetic
rint	rint(3M)	round to integral value in floating-point or integer format
isinf	ieee_functions(3M)	miscellaneous functions for IEEE arithmetic
isnan	ieee_functions(3M)	miscellaneous functions for IEEE arithmetic
isnormal	ieee_functions(3M)	miscellaneous functions for IEEE arithmetic
issubnormal	ieee_functions(3M)	miscellaneous functions for IEEE arithmetic
iszero	ieee_functions(3M)	miscellaneous functions for IEEE arithmetic
j0	bessel(3M)	Bessel functions
j1	bessel(3M)	Bessel functions
jn	bessel(3M)	Bessel functions
ldexp	frexp(3M)	traditional UNIX functions
lgamma	lgamma(3M)	log gamma function
log	exp(3M)	exponential, logarithm, power
log2	exp(3M)	exponential, logarithm, power
log10	exp(3M)	exponential, logarithm, power
log1p	exp(3M)	exponential, logarithm, power
logb	ieee_test(3M)	IEEE test functions for verifying standard compliance

**matherr**  
**max\_normal**  
**max\_subnormal**  
**min\_normal**  
**min\_subnormal**  
**modf**  
**nextafter**  
**nint**  
**pow**  
**quiet\_nan**  
**remainder**  
**rint**  
**scalb**  
**scalbn**  
**signaling\_nan**  
**signbit**  
**significant**  
**sin**  
**single\_precision**  
**sinh**  
**sqrt**  
**tan**  
**tanh**  
**y0**  
**y1**  
**yn**

**matherr(3M)**  
**ieee\_values(3M)**  
**ieee\_values(3M)**  
**ieee\_values(3M)**  
**ieee\_values(3M)**  
**frexp(3M)**  
**ieee\_functions(3M)**  
**rint(3M)**  
**exp(3M)**  
**ieee\_values(3M)**  
**ieee\_functions(3M)**  
**rint(3M)**  
**ieee\_test(3M)**  
**ieee\_functions(3M)**  
**ieee\_values(3M)**  
**ieee\_functions(3M)**  
**ieee\_test(3M)**  
**trig(3M)**  
**single\_precision(3M)**  
**hyperbolic(3M)**  
**sqrt(3M)**  
**trig(3M)**  
**hyperbolic(3M)**  
**bessel(3M)**  
**bessel(3M)**  
**bessel(3M)**

**math library exception-handling function**  
**functions that return extreme values of IEEE arithmetic**  
**functions that return extreme values of IEEE arithmetic**  
**functions that return extreme values of IEEE arithmetic**  
**functions that return extreme values of IEEE arithmetic**  
**traditional UNIX functions**  
**miscellaneous functions for IEEE arithmetic**  
**round to integral value in floating-point or integer format**  
**exponential, logarithm, power**  
**functions that return extreme values of IEEE arithmetic**  
**miscellaneous functions for IEEE arithmetic**  
**round to integral value in floating-point or integer format**  
**IEEE test functions for verifying standard compliance**  
**miscellaneous functions for IEEE arithmetic**  
**functions that return extreme values of IEEE arithmetic**  
**miscellaneous functions for IEEE arithmetic**  
**IEEE test functions for verifying standard compliance**  
**trigonometric functions**  
**single-precision access to libm functions**  
**hyperbolic functions**  
**cube root, square root**  
**trigonometric functions**  
**hyperbolic functions**  
**Bessel functions**  
**Bessel functions**  
**Bessel functions**

**NAME**

$j_0, j_1, j_n, y_0, y_1, y_n$  – Bessel functions

**SYNOPSIS**

```
#include <math.h>

double j0(x)
double x;

double j1(x)
double x;

double jn(n, x)
double x;
int n;

double y0(x)
double x;

double y1(x)
double x;

double yn(n, x)
double x;
int n;
```

**DESCRIPTION**

These functions calculate Bessel functions of the first and second kinds for real arguments and integer orders.

**SEE ALSO**

**exp(3M)**

**DIAGNOSTICS**

The functions  $y_0$ ,  $y_1$ , and  $y_n$  have logarithmic singularities at the origin, so they treat zero and negative arguments the way *log* does, as described in **exp(3M)**. Such arguments are unexceptional for  $j_0$ ,  $j_1$ , and  $j_n$ .



**NAME**

erf, erfc – error functions

**SYNOPSIS****#include <math.h>****double erf(x)****double x;****double erfc(x)****double x;****DESCRIPTION****erf(x)** returns the error function of  $x$ ; where  $\text{erf}(x) := (2/\sqrt{\pi}) \int_0^x \exp(-t^2) dt$ .**erfc(x)** returns  $1.0 - \text{erf}(x)$ , computed however by other methods that avoid cancellation for large  $x$ .

## NAME

exp, expm1, exp2, exp10, log, log1p, log2, log10, pow, compound, annuity – exponential, logarithm, power

## SYNOPSIS

```
#include <math.h>

double exp(x)
double x;

double expm1(x)
double x;

double exp2(x)
double x;

double exp10(x)
double x;

double log(x)
double x;

double log1p(x)
double x;

double log2(x)
double x;

double log10(x)
double x;

double pow(x, y)
double x, y;

double compound(r, n)
double r, n;

double annuity(r, n)
double r, n;
```

## DESCRIPTION

**exp()** returns the exponential function  $e^{**x}$ .

**expm1()** returns  $e^{**x}-1$  accurately even for tiny  $x$ .

**exp2()** and **exp10()** return  $2^{**x}$  and  $10^{**x}$  respectively.

**log()** returns the natural logarithm of  $x$ .

**log1p()** returns  $\log(1+x)$  accurately even for tiny  $x$ .

**log2()** and **log10()** return the logarithm to base 2 and 10 respectively.

**pow()** returns  $x^{**y}$ . **pow(x, 0.0)** is 1 for all  $x$ , in conformance with 4.3BSD, as discussed in the *Numerical Computation Guide*.

**compound()** and **annuity()** are functions important in financial computations of the effect of interest at periodic rate  $r$  over  $n$  periods. **compound( $r, n$ )** computes  $(1+r)^{**n}$ , the compound interest factor. Given an initial principal  $P0$ , its value after  $n$  periods is just  $Pn = P0 * \text{compound}(r, n)$ . **annuity( $r, n$ )** computes  $(1 - (1+r)^{**-n})/r$ , the present value of annuity factor. Given an initial principal  $P0$ , the equivalent periodic payment is just  $p = P0 / \text{annuity}(r, n)$ . **compound()** and **annuity()** are computed using **log1p()** and **expm1()** to avoid gratuitous inaccuracy for small-magnitude  $r$ . **compound()** and **annuity()** are not defined for  $r \leq -1$ .

Thus a principal amount  $P0$  placed at 5% annual interest compounded quarterly for 30 years would yield

$$P30 = P0 * \text{compound}(.05/4, 30.0 * 4)$$

while a conventional fixed-rate 30-year home loan of amount  $P0$  at 10% annual interest would be amortized by monthly payments in the amount

$$p = P0 / \text{annuity}(.10/12, 30.0 * 12)$$

SEE ALSO

**matherr(3M)**

DIAGNOSTICS

All these functions handle exceptional arguments in the spirit of ANSI/IEEE Std 754-1985. Thus for  $x == \pm 0$ ,  $\log(x)$  is  $-\infty$  with a division by zero exception; for  $x < 0$ , including  $-\infty$ ,  $\log(x)$  is a quiet NaN with an invalid operation exception; for  $x == +\infty$  or a quiet NaN,  $\log(x)$  is  $x$  without exception; for  $x$  a signaling NaN,  $\log(x)$  is a quiet NaN with an invalid operation exception; for  $x == 1$ ,  $\log(x)$  is 0 without exception; for any other positive  $x$ ,  $\log(x)$  is a normalized number with an inexact exception.

In addition,  $\exp()$ ,  $\exp2()$ ,  $\exp10()$ ,  $\log()$ ,  $\log2()$ ,  $\log10()$  and  $\text{pow}()$  may also set **errno** and call **matherr(3M)**.

## NAME

frexp, modf, ldexp – traditional UNIX functions

## SYNOPSIS

```
#include <math.h>
```

```
double frexp(value, eptr)
```

```
double value;
```

```
int *eptr;
```

```
double ldexp(x,n)
```

```
double x;
```

```
int n;
```

```
double modf(value, iptr)
```

```
double value, *iptr;
```

## DESCRIPTION

These functions are provided for compatibility with other UNIX system implementations. They are not used internally in `libm` or `libc`. Better ways to accomplish similar ends may be found in `ieee_functions(3M)` and `rint(3M)`.

`ldexp(x,n)` returns  $x * 2^{**n}$  computed by exponent manipulation rather than by actually performing an exponentiation or a multiplication. Note: `ldexp(x,n)` differs from `scalbn(x,n)`, defined in `ieee_functions(3M)`, only that in the event of IEEE overflow and underflow, `ldexp(x,n)` sets `errno` to `ERANGE`.

Every non-zero number can be written uniquely as  $x * 2^{**n}$ , where the significant  $x$  is in the range  $0.5 \leq |x| < 1.0$  and the exponent  $n$  is an integer. The function `frexp()` returns the significant of a double *value* as a double quantity,  $x$ , and stores the exponent  $n$ , indirectly through *eptr*. If *value* == 0, both results returned by `frexp()` are 0.

`modf()` returns the fractional part of *value* and stores the integral part indirectly through *iptr*. Thus the argument *value* and the returned values `modf()` and *iptr* satisfy

$$(*iptr + modf) == value$$

and both results have the same sign as *value*. The definition of `modf()` varies among UNIX system implementations, so avoid `modf()` in portable code.

The results of `frexp()` and `modf()` are not defined when *value* is an IEEE infinity or NaN.

## SEE ALSO

`ieee_functions(3M)`, `rint(3M)`

**NAME**

`sinh`, `cosh`, `tanh`, `asinh`, `acosh`, `atanh` – hyperbolic functions

**SYNOPSIS**

```
#include <math.h>

double sinh(x)
double x;

double cosh(x)
double x;

double tanh(x)
double x;

double asinh(x)
double x;

double acosh(x)
double x;

double atanh(x)
double x;
```

**DESCRIPTION**

These functions compute the designated direct and inverse hyperbolic functions for real arguments. They inherit much of their roundoff error from `expm1()` and `log1p`, described in `exp(3M)`.

**DIAGNOSTICS**

These functions handle exceptional arguments in the spirit of ANSI/IEEE Std 754-1985. Thus `sinh()` and `cosh()` return  $\pm\infty$  on overflow, `acosh()` returns a NaN if its argument is less than 1, and `atanh()` returns a NaN if its argument has absolute value greater than 1. In addition, `sinh`, `cosh`, and `tanh()` may also set `errno` and call `matherr(3M)`.

**SEE ALSO**

`exp(3M)`, `matherr(3M)`

## NAME

hypot – Euclidean distance

## SYNOPSIS

```
#include <math.h>
```

```
double hypot(x, y)
```

```
double x, y;
```

## DESCRIPTION

**hypot()** returns

```
sqrt(x*x + y*y),
```

taking precautions against unwarranted IEEE exceptions. On IEEE overflow, **hypot()** may also set **errno** and call **matherr(3M)**. **hypot( $\pm\infty$ , y)** is  $+\infty$  for any y, even a NaN, and is exceptional only for a signaling NaN.

**hypot(x,y)** and **atan2(y,x)** (see **trig(3M)**) convert rectangular coordinates (x,y) to polar (r, $\theta$ ); **hypot()** computes r, the modulus or radius.

## SEE ALSO

**trig(3M)**, **matherr(3M)**

**NAME**

`ieee_flags` – mode and status function for IEEE standard arithmetic

**SYNOPSIS**

```
#include <sys/ieeefp.h>
```

```
int ieee_flags(action, mode, in, out)
```

```
char *action, *mode, *in, **out;
```

**DESCRIPTION**

This function provides easy access to the modes and status required to fully exploit ANSI/IEEE Std 754-1985 arithmetic in a C program. All arguments are pointers to strings. Results arising from invalid arguments and invalid combinations are undefined for efficiency.

There are four types of *action*: **get**, **set**, **clear** and **clearall**. There are three valid settings for *mode*, two corresponding to modes of IEEE arithmetic:

**direction**            current rounding direction mode

**precision**           current rounding precision mode

and one corresponding to status of IEEE arithmetic:

**exception**            accrued exception-occurred status

There are fourteen types of *in* and *out*:

**nearest**                round toward nearest

**tozero**                round toward zero

**negative**              round toward negative infinity

**positive**              round toward positive infinity

**extended**

**double**

**single**

**inexact**

**division**              division by zero exception

**underflow**

**overflow**

**invalid**

**all**                    all five exceptions above

**common**                invalid, overflow, and division exceptions

Note: **all** and **common** only make sense with **set** or **clear**.

For **clearall**, `ieee_flags()` returns 0 and restores all default modes and status. Nothing will be assigned to *out*. Thus

```
char *mode, *out, *in;
```

```
ieee_flags("clearall", mode, in, &out);
```

set rounding direction to **nearest**, rounding precision to **extended**, and all accrued exception-occurred status to zero.

For `clear`, `ieee_flags()` returns 0 and restores the default mode or status. Nothing will be assigned to `out`. Thus

```
char *out, *in;
ieee_flags("clear", "direction", in, &out);    ... set rounding direction to round to nearest.
```

For `set`, `ieee_flags()` returns 0 if the action is successful and 1 if the corresponding required status or mode is not available (for instance, not supported in hardware). Nothing will be assigned to `out`. Thus

```
char *out, *in;
ieee_flags("set", "direction", "tozero", &out);    set rounding direction to round toward zero;
```

For `get`, we have the following cases:

Case 1: `mode` is `direction`. In that case, `out` returns one of the four strings `nearest`, `tozero`, `positive`, `negative`, and `ieee_flags()` returns a value corresponding to `out` according to the `enum fp_direction_type` defined in `<sys/ieeefp.h>`.

Case 2: `mode` is `precision`. In that case, `out` returns one of the three strings `extended`, `double` and `single`, and `ieee_flags()` returns a value corresponding to `out` according to the `enum fp_precision_type` defined in `<sys/ieeefp.h>`.

Case 3: `mode` is `exception`. In that case, `out` returns

```
not available    if information on exception is not available.
no exception     if no accrued exception.
```

the accrued exception that has the highest priority according to the following list:

```
the exception named by in
invalid
overflow
division
underflow
inexact
```

In this case `ieee_flags()` returns a five or six bit value where each bit (see `enum fp_exception_type` in `<sys/ieeefp.h>`) corresponds to an exception-occurred accrued status flag: 0 = off, 1 = on. The bit corresponding to a particular exception varies among architectures (see `<sys/ieeefp.h>`).

Example:

```
char *out; int k, ieee_flags();
ieee_flags("clear", "exception", "all", &out);    /* clear all accrued exceptions */
...
code that generates three exceptions: overflow, invalid, inexact
...
k = ieee_flags("get", "exception", "overflow", &out);
```

then `out` is `overflow`, and on a Sun-3, `k` is 25.



**NAME**

ieee\_functions, fp\_class, finite, ilogb, isinf, isnan, isnormal, issubnormal, iszero, signbit, copysign, fabs, fmod, nextafter, remainder, scalbn – appendix and related miscellaneous functions for IEEE arithmetic

**SYNOPSIS**

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

enum fp_class_type fp_class(x)
double x;

int finite(x)
double x;

int ilogb(x)
double x;

int isinf(x)
double x;

int isnan(x)
double x;

int isnormal(x)
double x;

int issubnormal(x)
double x;

int iszero(x)
double x;

int signbit(x)
double x;

void ieee_retrospective(f)
FILE *f;

void nonstandard_arithmetic()
void standard_arithmetic()

double copysign(x,y)
double x, y;

double fabs(x)
double x;

double fmod(x,y)
double x, y;

double nextafter(x,y)
double x, y;

double remainder(x,y)
double x, y;

double scalbn(x,n)
double x; int n;
```

**DESCRIPTION**

Most of these functions provide capabilities required by ANSI/IEEE Std 754-1985 or suggested in its appendix.

**fp\_class(x)** corresponds to the IEEE's **class()** and classifies  $x$  as zero, subnormal, normal,  $\infty$ , or quiet or signaling *NaN*. `<floatingpoint.h>` defines **enum fp\_class\_type**. The following functions return 0 if the indicated condition is not satisfied:

<b>finite(x)</b>	returns 1 if $x$ is zero, subnormal or normal
<b>isinf(x)</b>	returns 1 if $x$ is $\infty$
<b>isnan(x)</b>	returns 1 if $x$ is <i>NaN</i>
<b>isnormal(x)</b>	returns 1 if $x$ is normal
<b>issubnormal(x)</b>	returns 1 if $x$ is subnormal
<b>iszero(x)</b>	returns 1 if $x$ is zero
<b>signbit(x)</b>	returns 1 if $x$ 's sign bit is set

**ilogb(x)** returns the unbiased exponent of  $x$  in integer format. **ilogb( $\pm\infty$ )** = +MAXINT and **ilogb(0)** = -MAXINT; `<values.h>` defines MAXINT as the largest int. **ilogb(x)** never generates an exception. When  $x$  is subnormal, **ilogb(x)** returns an exponent computed as if  $x$  were first normalized.

**ieee\_retrospective(f)** prints a message to the FILE  $f$  listing all IEEE accrued exception-occurred bits currently on, unless no such bits are on or the only one on is "inexact". It's intended to be used at the end of a program to indicate whether some IEEE floating-point exceptions occurred that might have affected the result.

**standard\_arithmetic()** and **nonstandard\_arithmetic()** are meaningful on systems that provide an alternative faster mode of floating-point arithmetic that does not conform to the default IEEE Standard. Nonstandard modes vary among implementations; nonstandard mode may, for instance, result in setting subnormal results to zero or in treating subnormal operands as zero, or both, or something else. **standard\_arithmetic()** reverts to the default standard mode. On systems that provide only one mode, these functions have no effect.

**copysign(x,y)** returns  $x$  with  $y$ 's sign bit.

**fabs(x)** returns the absolute value of  $x$ .

**nextafter(x,y)** returns the next machine representable number from  $x$  in the direction  $y$ .

**remainder(x, y)** and **fmod(x, y)** return a remainder of  $x$  with respect to  $y$ ; that is, the result  $r$  is one of the numbers that differ from  $x$  by an integral multiple of  $y$ . Thus  $(x - r)/y$  is an integral value, even though it might exceed MAXINT if it were explicitly computed as an int. Both functions return one of the two such  $r$  smallest in magnitude. **remainder(x, y)** is the operation specified in ANSI/IEEE Std 754-1985; the result of **fmod(x, y)** may differ from **remainder()**'s result by  $\pm y$ . The magnitude of **remainder()**'s result can not exceed half that of  $y$ ; its sign might not agree with either  $x$  or  $y$ . The magnitude of **fmod()**'s result is less than that of  $y$ ; its sign agrees with that of  $x$ . Neither function can generate an exception as long as both arguments are normal or subnormal. **remainder(x, 0)**, **fmod(x, 0)**, **remainder( $\infty$ , y)**, and **fmod( $\infty$ , y)** are invalid operations that produce a *NaN*.

**scalbn(x, n)** returns  $x * 2^{**n}$  computed by exponent manipulation rather than by actually performing an exponentiation or a multiplication. Thus

$$1 \leq \text{scalbn}(\text{fabs}(x), -\text{ilogb}(x)) < 2$$

for every  $x$  except 0,  $\infty$ , and *NaN*.

**SEE ALSO**

**floatingpoint(3)**, **ieee\_flags(3M)**, **matherr(3M)**

**NAME**

`ieee_handler` – IEEE exception trap handler function

**SYNOPSIS**

```
#include <floatingpoint.h>

int ieee_handler(action,exception,hdl)
char action[ ], exception[ ];
sigfpe_handler_type hdl;
```

**DESCRIPTION**

This function provides easy exception handling to exploit ANSI/IEEE Std 754-1985 arithmetic in a C program. The first two arguments are pointers to strings. Results arising from invalid arguments and invalid combinations are undefined for efficiency.

There are three types of *action* : **get**, **set**, and **clear**. There are five types of *exception* :

<b>inexact</b>	
<b>division</b>	... division by zero exception
<b>underflow</b>	
<b>overflow</b>	
<b>invalid</b>	
<b>all</b>	... all five exceptions above
<b>common</b>	... invalid, overflow, and division exceptions

Note: **all** and **common** only make sense with **set** or **clear**.

**hdl** contains the address of a signal-handling routine. `<floatingpoint.h>` defines `sigfpe_handler_type`.

**get** will return the location of the current handler routine for *exception* cast to an int. **set** will set the routine pointed at by **hdl** to be the handler routine and at the same time enable the trap on *exception*, except when **hdl** == `SIGFPE_DEFAULT` or `SIGFPE_IGNORE`; then `ieee_handler()` will disable the trap on *exception*. When **hdl** == `SIGFPE_ABORT`, any trap on *exception* will dump core using `abort(3)`. **clear** all disables trapping on all five exceptions.

Two steps are required to intercept an IEEE-related SIGFPE code with `ieee_handler`:

- 1) Set up a handler with `ieee_handler`.
- 2) Perform a floating-point operation that generates the intended IEEE exception.

Unlike `sigfpe(3)`, `ieee_handler()` also adjusts floating-point hardware mode bits affecting IEEE trapping. For **clear**, **set** `SIGFPE_DEFAULT`, or **set** `SIGFPE_IGNORE`, the hardware trap is disabled. For any other **set**, the hardware trap is enabled.

SIGFPE signals can be handled using `sigvec(2)`, `signal(3V)`, `sigfpe(3)`, or `ieee_handler(3M)`. In a particular program, to avoid confusion, use only one of these interfaces to handle SIGFPE signals.

**DIAGNOSTICS**

`ieee_handler()` normally returns 0 for **set**. 1 will be returned if the action is not available (for instance, not supported in hardware). For **get**, the address of the current handler is returned, cast to an int.

**EXAMPLE**

A user-specified signal handler might look like this:

```
void sample_handler(sig, code, scp, addr)
int sig;      /* sig == SIGFPE always */
int code;
struct sigcontext *scp;
char *addr;
{
    /*
     * Sample user-written sigfpe code handler.
     * Prints a message and continues.
     * struct sigcontext is defined in <signal.h>.
     */
    printf("ieee exception code %x occurred at pc %X \n", code, scp->sc_pc);
}
```

and it might be set up like this:

```
extern void sample_handler();
main()
{
    sigfpe_handler_type hdl, old_handler1, old_handler2;
    /*
     * save current overflow and invalid handlers
     */
    old_handler1 = (sigfpe_handler_type) ieee_handler("get", "overflow", old_handler1);
    old_handler2 = (sigfpe_handler_type) ieee_handler("get", "invalid", old_handler2);
    /*
     * set new overflow handler to sample_handler() and set new
     * invalid handler to SIGFPE_ABORT (abort on invalid)
     */
    hdl = (sigfpe_handler_type) sample_handler;
    if (ieee_handler("set", "overflow", hdl) != 0)
        printf("ieee_handler can't set overflow \n");
    if (ieee_handler("set", "invalid", SIGFPE_ABORT) != 0)
        printf("ieee_handler can't set invalid \n");
    ...
    /*
     * restore old overflow and invalid handlers
     */
    ieee_handler("set", "overflow", old_handler1);
    ieee_handler("set", "invalid", old_handler2);
}
```

**SEE ALSO**

sigvec(2), abort(3), floatingpoint(3), sigfpe(3), signal(3V)

**NAME**

ieee\_test, logb, scalb, significant – IEEE test functions for verifying standard compliance

**SYNOPSIS**

```
#include <math.h>

double logb(x)
double x;

double scalb(x,y)
double x; double y;

double significant(x)
double x;
```

**DESCRIPTION**

These functions allow users to verify compliance to ANSI/IEEE Std 754-1985 by running certain test vectors distributed by the University of California. Their use is not otherwise recommended; instead use `scalbn(x,n)` and `ilogb(x)` described in `ieee_functions(3M)`. See the *Numerical Computation Guide* for details.

`logb(x)` returns the unbiased exponent of  $x$  in floating-point format, for exercising the `logb(L)` test vector. `logb( $\pm\infty$ ) =  $+\infty$` ; `logb(0) =  $-\infty$`  with a division by zero exception. `logb(x)` differs from `ilogb(x)` in returning a result in floating-point rather than integer format, in sometimes signaling IEEE exceptions, and in not normalizing subnormal  $x$ .

`scalb(x,(double)n)` returns  $x * 2^{**n}$  computed by exponent manipulation rather than by actually performing an exponentiation or a multiplication, for exercising the `scalb(S)` test vector. Thus

$$0 \leq \text{scalb}(\text{fabs}(x), -\text{logb}(x)) < 2$$

for every  $x$  except 0,  $\infty$  and *NaN*. `scalb(x,y)` is not defined when  $y$  is not an integral value. `scalb(x,y)` differs from `scalbn(x,n)` in that the second argument is in floating-point rather than integer format.

`significant(x)` computes just

$$\text{scalb}(x, (\text{double}) -\text{ilogb}(x)),$$

for exercising the `fraction-part(F)` test vector.

**FILES**

/usr/lib/libm.a

**SEE ALSO**

`floatingpoint(3)`, `ieee_values(3M)`, `ieee_functions(3M)`, `matherr(3M)`

**NAME**

ieee\_values, min\_subnormal, max\_subnormal, min\_normal, max\_normal, infinity, quiet\_nan, signaling\_nan, HUGE, HUGE\_VAL – functions that return extreme values of IEEE arithmetic

**SYNOPSIS**

```
#include <math.h>

double min_subnormal()
double max_subnormal()
double min_normal()
double max_normal()
double infinity()
double quiet_nan(n)
long n;
double signaling_nan(n)
long n;
#define HUGE (infinity())
#define HUGE_VAL (infinity())
```

**DESCRIPTION**

These functions return special values associated with ANSI/IEEE Std 754-1985 double-precision floating-point arithmetic: the smallest and largest positive subnormal numbers, the smallest and largest positive normalized numbers, positive infinity, and a quiet and signaling NaN. The long parameters  $n$  to `quiet_nan( $n$ )` and `signaling_nan( $n$ )` are presently unused but are reserved for future use to specify the significant of the returned NaN.

None of these functions are affected by IEEE rounding or trapping modes or generate any IEEE exceptions.

The macro `HUGE` returns  $+\infty$  in accordance with previous SunOS releases. The macro `HUGE_VAL` returns  $+\infty$  in accordance with the System V Interface Definition.

**FILES**

/usr/lib/libm.a

**SEE ALSO**

ieee\_functions(3M)

**NAME**

lgamma – log gamma function

**SYNOPSIS**

```
#include <math.h>
extern int signgam;
double lgamma(x)
double x;
```

**DESCRIPTION**

lgamma() returns

$$\ln |\Gamma(x)|$$

where

$$\Gamma(x) = \int_0^{\infty} t^{x-1} e^{-t} dt$$

for  $x > 0$  and

$$\Gamma(x) = \pi / (\Gamma(1-x) \sin(\pi x))$$

for  $x < 1$ .

The external integer **signgam** returns the sign of  $\Gamma(x)$ .

**IDIOSYNCRASIES**

Do *not* use the expression **signgam\*exp(lgamma(x))** to compute '**g :=  $\Gamma(x)$** '. Instead compute **lgamma()** first:

```
lg = lgamma(x); g = signgam*exp(lg);
```

only after **lgamma()** has returned can **signgam** be correct. Note:  $\Gamma(x)$  must overflow when  $x$  is large enough, underflow when  $-x$  is large enough, and generate a division by zero exception at the singularities  $x$  a nonpositive integer. In addition, **lgamma()** may also set **errno** and call **matherr(3M)**.

**SEE ALSO****matherr(3M)**

## NAME

matherr – math library exception-handling function

## SYNOPSIS

```
#include <math.h>

int matherr(exc)
struct exception *exc;
```

## DESCRIPTION

The SVID (*System V Interface Definition*) specifies that certain **libm** functions call **matherr()** when exceptions are detected. Users may define their own mechanisms for handling exceptions, by including a function named **matherr()** in their programs. **matherr()** is of the form described above. When an exception occurs, a pointer to the exception structure *exc* will be passed to the user-supplied **matherr()** function. This structure, which is defined in the **<math.h>** header file, is as follows:

```
struct exception {
    int type;
    char *name;
    double arg1, arg2, retval;
};
```

The element **type** is an integer describing the type of exception that has occurred, from the following list of constants (defined in the header file):

<b>DOMAIN</b>	argument domain exception
<b>SING</b>	argument singularity
<b>OVERFLOW</b>	overflow range exception
<b>UNDERFLOW</b>	underflow range exception

The element **name** points to a string containing the name of the function that incurred the exception. The elements **arg1** and **arg2** are the arguments with which the function was invoked. **retval** is set to the default value that will be returned by the function unless the user's **matherr()** sets it to a different value.

If the user's **matherr()** function returns non-zero, no exception message will be printed, and **errno** will not be set.

If **matherr()** is not supplied by the user, the default **matherr** exception-handling mechanisms, summarized in the table below, will be invoked upon exception:

**DOMAIN==fp\_invalid**

An IEEE NaN is usually returned, **errno** is set to EDOM, and a message is printed on standard error. **pow(0.0,0.0)** and **atan2(0.0,0.0)** return numerical default results but set **errno** and print the message.

**SING==fp\_division**

An IEEE  $\infty$  of appropriate sign is returned, **errno** is set to EDOM, and a message is printed on standard error.

**OVERFLOW==fp\_overflow**

In the default rounding direction, an IEEE  $\infty$  of appropriate sign is returned. In optional rounding directions,  $\pm$ MAXDOUBLE, the largest finite double-precision number, is sometimes returned instead of  $\pm\infty$ . **errno** is set to ERANGE.

**UNDERFLOW==fp\_underflow**

An appropriately-signed zero, subnormal number, or smallest normalized number is returned, and **errno** is set to ERANGE.

The facilities provided by **matherr()** are not available in situations such as compiling on a Sun-3 system with **/usr/lib/f68881/libm.il** or **/usr/lib/ffpa/libm.il**, in which case some **libm** functions are converted to atomic hardware operations. In these cases setting **errno** and calling **matherr()** are not worth the adverse performance impact, but regular ANSI/IEEE Std 754-1985 exception handling remains available. In any



case **errno** is not a reliable error indicator in that it may be unexpectedly set by a function in a handler for an asynchronous signal.

DEFAULT ERROR HANDLING PROCEDURES				
<i>Types of Errors</i>				
<math.h> type	DOMAIN	SING	OVERFLOW	UNDERFLOW
<b>errno</b>	EDOM	EDOM	ERANGE	ERANGE
IEEE Exception	Invalid Operation	Division by Zero	Overflow	Underflow
<floatingpoint.h> type	fp_invalid	fp_division	fp_overflow	fp_underflow
ACOS, ASIN:	M, NaN	–	–	–
ATAN2(0,0):	M, $\pm 0.0$ or $\pm\pi$	–	–	–
BESSEL: y0, y1, yn ( $x < 0$ )	M, NaN	–	–	–
y0, y1, yn ( $x = 0$ )	–	M, $-\infty$	–	–
COSH, SINH:	–	–	IEEE Overflow	–
EXP:	–	–	IEEE Overflow	IEEE Underflow
HYPOT:	–	–	IEEE Overflow	–
LGAMMA:	–	M, $+\infty$	IEEE Overflow	–
LOG, LOG10: ( $x < 0$ )	M, NaN	–	–	–
( $x = 0$ )	–	M, $-\infty$	–	–
POW: usual cases	–	–	IEEE Overflow	IEEE Underflow
( $x < 0$ ) ** (y not an integer)	M, NaN	–	–	–
0 ** 0	M, 1.0	–	–	–
0 ** (y < 0)	–	M, $\pm\infty$	–	–
SQRT:	M, NaN	–	–	–

#### ABBREVIATIONS

M	Message is printed (EDOM exception).
NaN	IEEE NaN result and invalid operation exception.
$\infty$	IEEE $\infty$ result and division-by-zero exception.
IEEE Overflow	IEEE Overflow result and exception.
IEEE Underflow	IEEE Underflow result and exception.
$\pi$	Closest machine-representable approximation to pi.

The interaction of IEEE arithmetic and **matherr()** is not defined when executing under IEEE rounding modes other than the default round to nearest: **matherr()** may not be called on overflow or underflow, and the Sun-provided **matherr()** may return results that differ from those in this table.

**EXAMPLE**

```
#include <math.h>

int
matherr(x)
register struct exception *x;
{
    switch (x->type) {
    case
        DOMAIN:
            /* change sqrt to return sqrt(-arg1), not NaN */
            if (!strcmp(x->name, "sqrt")) {
                x->retval = sqrt(-x->arg1);
                return (0); /* print message and set errno */
            } /* fall through */
    case
        SING:
            /* all other domain or sing exceptions, print message and abort */
            fprintf(stderr, "domain exception in %s\n", x->name);
            abort();
            break;
    }
    return (0); /* all other exceptions, execute default procedure */
}
```

**NAME**

**aint**, **anint**, **ceil**, **floor**, **rint**, **irint**, **nint** – round to integral value in floating-point or integer format

**SYNOPSIS**

```
#include <math.h>
```

```
double aint(x)
```

```
double x;
```

```
double anint(x)
```

```
double x;
```

```
double ceil(x)
```

```
double x;
```

```
double floor(x)
```

```
double x;
```

```
double rint(x)
```

```
double x;
```

```
int irint(x)
```

```
double x;
```

```
int nint(x)
```

```
double x;
```

**DESCRIPTION**

**aint()**, **anint()**, **ceil()**, **floor()**, and **rint()** convert a double value into an integral value in double format. They vary in how they choose the result when the argument is not already an integral value. Here an “integral value” means a value of a mathematical integer, which however might be too large to fit in a particular computer’s int format. All sufficiently large values in a particular floating-point format are already integral; in IEEE double-precision format, that means all values  $\geq 2^{52}$ . Zeros, infinities, and quiet NaNs are treated as integral values by these functions, which always preserve their argument’s sign.

**aint()** returns the integral value between  $x$  and 0, nearest  $x$ . This corresponds to IEEE rounding toward zero and to the Fortran generic intrinsic function **aint()**.

**anint()** returns the nearest integral value to  $x$ , except halfway cases are rounded to the integral value larger in magnitude. This corresponds to the Fortran generic intrinsic function **anint()**.

**ceil()** returns the least integral value greater than or equal to  $x$ . This corresponds to IEEE rounding toward positive infinity.

**floor()** returns the greatest integral value less than or equal to  $x$ . This corresponds to IEEE rounding toward negative infinity.

**rint()** rounds  $x$  to an integral value according to the current IEEE rounding direction.

**irint()** converts  $x$  into int format according to the current IEEE rounding direction.

**nint()** converts  $x$  into int format rounding to the nearest int value, except halfway cases are rounded to the int value larger in magnitude. This corresponds to the Fortran generic intrinsic function **nint()**.

## NAME

single\_precision – single-precision access to libm functions

## SYNOPSIS

```
#include <math.h>
```

```

FLOATFUNCTIONTYPE r_acos_(x)
FLOATFUNCTIONTYPE r_acospi_(x)
FLOATFUNCTIONTYPE r_acosh_(x)
FLOATFUNCTIONTYPE r_aint_(x)
FLOATFUNCTIONTYPE r_anint_(x)
FLOATFUNCTIONTYPE r_annuity_(x)
FLOATFUNCTIONTYPE r_asin_(x)
FLOATFUNCTIONTYPE r_asinpi_(x)
FLOATFUNCTIONTYPE r_asinh_(x)
FLOATFUNCTIONTYPE r_atan_(x)
FLOATFUNCTIONTYPE r_atanpi_(x)
FLOATFUNCTIONTYPE r_atanh_(x)
FLOATFUNCTIONTYPE r_atan2_(x,y)
FLOATFUNCTIONTYPE r_atan2pi_(x,y)
FLOATFUNCTIONTYPE r_cbrt_(x)
FLOATFUNCTIONTYPE r_ceil_(x)
enum fp_class_type ir_fp_class_(x)
FLOATFUNCTIONTYPE r_compound_(x,y)
FLOATFUNCTIONTYPE r_copysign_(x,y)
FLOATFUNCTIONTYPE r_cos_(x)
FLOATFUNCTIONTYPE r_cospi_(x)
FLOATFUNCTIONTYPE r_cosh_(x)
FLOATFUNCTIONTYPE r_erf_(x)
FLOATFUNCTIONTYPE r_erfc_(x)
FLOATFUNCTIONTYPE r_exp_(x)
FLOATFUNCTIONTYPE r_expm1_(x)
FLOATFUNCTIONTYPE r_exp2_(x)
FLOATFUNCTIONTYPE r_exp10_(x)
FLOATFUNCTIONTYPE r_fabs_(x)
int ir_finite_(x)
FLOATFUNCTIONTYPE r_floor_(x)
FLOATFUNCTIONTYPE r_fmod_(x,y)
FLOATFUNCTIONTYPE r_hypot_(x,y)
int ir_ilogb_(x)
int ir_rint_(x)
int ir_isinf_(x)
int ir_isnan_(x)
int ir_isnormal_(x)
int ir_issubnormal_(x)
int ir_iszero_(x)
int ir_nint_(x)
FLOATFUNCTIONTYPE r_infinity_( )
FLOATFUNCTIONTYPE r_j0_(x)
FLOATFUNCTIONTYPE r_j1_(x)
FLOATFUNCTIONTYPE r_jn_(n,x)
FLOATFUNCTIONTYPE r_lgamma_(x)
FLOATFUNCTIONTYPE r_logb_(x)
FLOATFUNCTIONTYPE r_log_(x)
FLOATFUNCTIONTYPE r_log1p_(x)

```

```

FLOATFUNCTIONTYPE r_log2_(x)
FLOATFUNCTIONTYPE r_log10_(x)
FLOATFUNCTIONTYPE r_max_normal_( )
FLOATFUNCTIONTYPE r_max_subnormal_( )
FLOATFUNCTIONTYPE r_min_normal_( )
FLOATFUNCTIONTYPE r_min_subnormal_( )
FLOATFUNCTIONTYPE r_nextafter_(x,y)
FLOATFUNCTIONTYPE r_pow_(x,y)
FLOATFUNCTIONTYPE r_quiet_nan_(n)
FLOATFUNCTIONTYPE r_remainder_(x,y)
FLOATFUNCTIONTYPE r_rint_(x)
FLOATFUNCTIONTYPE r_scalb_(x,y)
FLOATFUNCTIONTYPE r_scalbn_(x,n)
FLOATFUNCTIONTYPE r_signaling_nan_(n)
int ir_signbit_(x)
FLOATFUNCTIONTYPE r_significant_(x)
FLOATFUNCTIONTYPE r_sin_(x)
FLOATFUNCTIONTYPE r_sinpi_(x)
void r_sincos_(x,s,c)
void r_sincospi_(x,s,c)
FLOATFUNCTIONTYPE r_sinh_(x)
FLOATFUNCTIONTYPE r_sqrt_(x)
FLOATFUNCTIONTYPE r_tan_(x)
FLOATFUNCTIONTYPE r_tanpi_(x)
FLOATFUNCTIONTYPE r_tanh_(x)
FLOATFUNCTIONTYPE r_y0_(x)
FLOATFUNCTIONTYPE r_y1_(x)
FLOATFUNCTIONTYPE r_yn_(n,x)

float *x, *y, *s, *c
int *n

```

**DESCRIPTION**

These functions are single-precision versions of certain **libm** functions. Primarily for use by Fortran programmers, these functions may also be used in other languages. The single-precision floating-point results are deviously declared to avoid C's automatic type conversion to double.

**FILES**

/usr/lib/libm.a

**NAME**

sqrt, cbrt – cube root, square root

**SYNOPSIS**

```
#include <math.h>
```

```
double cbrt(x)
```

```
double x;
```

```
double sqrt(x)
```

```
double x;
```

**DESCRIPTION**

**sqrt(x)** returns the square root of  $x$ , correctly rounded according to ANSI/IEEE 754-1985. In addition, **sqrt()** may also set **errno** and call **matherr(3M)**.

**cbrt(x)** returns the cube root of  $x$ . **cbrt()** is accurate to within 0.7 *ulps*.

**SEE ALSO**

**matherr(3M)**

## NAME

sin, cos, tan, asin, acos, atan, atan2 – trigonometric functions

## SYNOPSIS

```
#include <math.h>

double sin(x)
double x;

double cos(x)
double x;

void sincos(x, s, c)
double x, *s, *c;

double tan(x)
double x;

double asin(x)
double x;

double acos(x)
double x;

double atan(x)
double x;

double atan2(y, x)
double y, x;

double sinpi(x)
double x;

double cospi(x)
double x;

void sincospi(x, s, c)
double x, *s, *c;

double tanpi(x)
double x;

double asinpi(x)
double x;

double acospi(x)
double x;

double atanpi(x)
double x;

double atan2pi(y, x)
double y, x;
```

## DESCRIPTION

`sin()`, `cos()`, `sincos()`, and `tan()` return trigonometric functions of radian arguments. The values of trigonometric functions of arguments exceeding  $\pi/4$  in magnitude are affected by the precision of the approximation to  $\pi/2$  used to reduce those arguments to the range  $-\pi/4$  to  $\pi/4$ . Argument reduction may occur in hardware or software; if in software, the variable `fp_pi` defined in `<math.h>` allows changing that precision at run time. Trigonometric argument reduction is discussed in the *Numerical Computation Guide*. Note: `sincos(x,s,c)` allows simultaneous computation of `*s = sin(x)` and `*c = cos(x)`.

`asin()` returns the arc-sin in the range  $-\pi/2$  to  $\pi/2$ .

**acos()** returns the arc cosine in the range 0 to  $\pi$ .

**atan()** returns the arc tangent of  $x$  in the range  $-\pi/2$  to  $\pi/2$ .

**atan2(y,x)** and **hypot(x,y)** (see **hypot(3M)**) convert rectangular coordinates  $(x,y)$  to polar  $(r,\theta)$ ; **atan2()** computes  $\theta$ , the argument or phase, by computing an arc tangent of  $y/x$  in the range  $-\pi$  to  $\pi$ . **atan2(0.0,0.0)** is  $\pm 0.0$  or  $\pm\pi$ , in conformance with 4.3BSD, as discussed in the *Numerical Computation Guide*.

**sinpi()**, **cospi()**, and **tanpi()** avoid range-reduction issues because their definition **sinpi(x)==sin( $\pi*x$ )** permits range reduction that is fast and exact for all  $x$ . The corresponding inverse functions compute **asinpi(x)==asin(x)/ $\pi$** . Similarly **atan2pi(y,x)==atan2(y,x)/ $\pi$** .

#### DIAGNOSTICS

These functions handle exceptional arguments in the spirit of ANSI/IEEE Std 754-1985. **sin( $\pm\infty$ )**, **cos( $\pm\infty$ )**, **tan( $\pm\infty$ )**, or **asin(x)** or **acos(x)** with  $|x|>1$ , return NaN; **sinpi(x)** et. al. are similar. In addition, **asin()**, **acos()**, and **atan2()** may also set **errno** and call **matherr(3M)**.

#### SEE ALSO

**hypot(3M)**, **matherr(3M)**



**NAME**

intro – introduction to RPC service library functions and protocols

**DESCRIPTION**

These functions constitute the RPC service library. Most of these describe RPC protocols. The PROTOCOL section describes how to access the protocol description file. This file may be compiled with `rpcgen(1)` to produce data definitions and XDR routines. Precompiled versions of header files sometimes exist as `<rpcsvc/*.h>` and precompiled XDR routines and programming interfaces to the protocols sometimes exist in `librpcsvc`. Warning: some of these header files and XDR routines were hand-written because they existed before `rpcgen`. They do not correspond to their protocol description file. In order to get the link editor to load this library, use the `-lrpcsvc` option of `cc(1V)`. Information about the availability of programming interfaces to these protocols is available under PROGRAMMING section of each manual page.

Some routines in the `librpcsvc` library do not correspond to protocols, but are useful utilities for RPC programming. These are distinguished by the presence of the SYNOPSIS section instead of the usual PROTOCOL section.

**LIST OF STANDARD RPC SERVICES**

Name	Appears on Page	Description
<code>bootparam</code>	<code>bootparam(3R)</code>	bootparam protocol
<code>ether</code>	<code>ether(3R)</code>	monitor traffic on the Ethernet
<code>getpublickey</code>	<code>publickey(3R)</code>	get public or secret key
<code>getrpcport</code>	<code>getrpcport(3R)</code>	get RPC port number
<code>getsecretkey</code>	<code>publickey(3R)</code>	get public or secret key
<code>ipalloc</code>	<code>ipalloc(3R)</code>	determine or temporarily allocate IP address
<code>klm_prot</code>	<code>klm_prot(3R)</code>	protocol between kernel and local lock manager
<code>mount</code>	<code>mount(3R)</code>	keep track of remotely mounted filesystems
<code>nlm_prot</code>	<code>nlm_prot(3R)</code>	protocol between local and remote network lock managers
<code>passwd2des</code>	<code>xcrypt(3R)</code>	hex encryption and utility routines
<code>pnp</code>	<code>pnp(3R)</code>	automatic network installation
<code>publickey</code>	<code>publickey(3R)</code>	get public or secret key
<code>rex</code>	<code>rex(3R)</code>	remote execution protocol
<code>rnusers</code>	<code>rnusers(3R)</code>	return information about users on remote machines
<code>rquota</code>	<code>rquota(3R)</code>	implement quotas on remote machines
<code>rstat</code>	<code>rstat(3R)</code>	get performance data from remote kernel
<code>rusers</code>	<code>rnusers(3R)</code>	return information about users on remote machines
<code>rwall</code>	<code>rwall(3R)</code>	write to specified remote machines
<code>sm_inter</code>	<code>sm_inter(3R)</code>	status monitor protocol
<code>spray</code>	<code>spray(3R)</code>	scatter data in order to check the network
<code>xcrypt</code>	<code>xcrypt(3R)</code>	hex encryption and utility routines
<code>xdecrypt</code>	<code>xcrypt(3R)</code>	hex encryption and utility routines
<code>xencrypt</code>	<code>xcrypt(3R)</code>	hex encryption and utility routines
<code>yp</code>	<code>yp(3R)</code>	NIS protocol
<code>yppasswd</code>	<code>yppasswd(3R)</code>	update user password in NIS

**NAME**

bootparam – bootparam protocol

**PROTOCOL**

/usr/include/rpcsvc/bootparam\_prot.x

**DESCRIPTION**

The bootparam protocol is used for providing information to the diskless clients necessary for booting.

**PROGRAMMING**

#include <rpcsvc/bootparam.h>

**XDR Routines**

The following XDR routines are available in **librpcsvc**:

xdr\_bp\_whoami\_arg

xdr\_bp\_whoami\_res

xdr\_bp\_getfile\_arg

xdr\_bp\_getfile\_res

**SEE ALSO**

bootparams(5), bootparamd(8)

**NAME**

ether – monitor traffic on the Ethernet

**PROTOCOL**

`/usr/include/rpcsvc/ether.x`

**DESCRIPTION**

The ether protocol is used for monitoring traffic on the ethernet.

**PROGRAMMING**

`#include <rpcsvc/ether.h>`

The following XDR routines are available in `librpcsvc`:

`xdr_etherstat`  
`xdr_etheraddr`  
`xdr_etherhtable`  
`xdr_etherhmem`  
`xdr_addrmask`

**SEE ALSO**

`traffic(1C)`, `etherfind(8C)`, `etherd(8C)`

**NAME**

getrpcport – get RPC port number

**SYNOPSIS**

```
int getrpcport(host, prognum, versnum, proto)  
char *host;  
int prognum, versnum, proto;
```

**DESCRIPTION**

**getrpcport()** returns the port number for version *versnum* of the RPC program *prognum* running on *host* and using protocol *proto*. It returns 0 if it cannot contact the portmapper, or if *prognum* is not registered. If *prognum* is registered but not with version *versnum*, it will still return a port number (for some version of the program) indicating that the program is indeed registered. The version mismatch will be detected upon the first call to the service.

**NAME**

`ipalloc` – determine or temporarily allocate IP address

**PROTOCOL**

`/usr/include/rpcsvc/ipalloc.x`

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

`ipalloc()` is the protocol for allocating the IP address that a system should use.

**PROGRAMMING**

```
#include <rpcsvc/ipalloc.h>
```

The following RPC calls are available in version 2 of this protocol:

**NULLPROC**

This is a standard null entry, used to ping a service to measure overhead or to discover servers.

**IP\_ALLOC**

Returns an IP address corresponding to a given Ethernet address, if possible. This RPC must be called using DES authentication, from a client authorized to allocate IP addresses. A cache of allocated addresses is maintained.

The first action taken on receipt of this RPC is to verify that no existing mapping between the *etheraddr* and the *netnum* exists in the Network Information Service (NIS) database. If one is found, then that is returned. Otherwise, an internal cache is checked, and if an entry is found there for the given *etheraddr* on the right network, that entry is used. If no address was found either in the NIS database or in the cache, a new one may be allocated and returned, and the *ip\_success* status is returned.

If an unusable entry was found in the cache, this RPC returns *ip\_failure* status.

**IP\_TONAME**

Used to determine whether a given IP address is known to the NIS service, since NIS allows a delay between the posting of an address and its availability in some locations on the network.

**IP\_FREE**

This RPC is used to delete *ipaddr* entries from the cache when they are no longer needed there. It requires the same protections as the **IP\_ALLOC** RPC.

**SEE ALSO**

`ipallocald(8C)`, `pnpboot(8C)`

**NOTES**

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed.

**NAME**

klm\_prot – protocol between kernel and local lock manager

**PROTOCOL**

/usr/include/klm\_prot.x

**DESCRIPTION**

The protocol is used for communication between kernel and local lock manager.

**PROGRAMMING**

```
#include <rpcsvc/klm_prot.h>
```

**XDR Routines**

The following XDR routines are available in **librpcsvc**:

```
xdr_klm_testargs  
xdr_klm_testrply  
xdr_klm_lockargs  
xdr_klm_unlockargs  
xdr_klm_stat
```

**SEE ALSO**

lockd(8C)

**NAME**

mount – keep track of remotely mounted filesystems

**PROTOCOL**

**/usr/include/rpcsvc/mount.x**

**DESCRIPTION**

The mount protocol is separate from, but related to, the NFS protocol. It provides all of the operating system specific services to get the NFS off the ground — looking up path names, validating user identity, and checking access permissions. Clients use the mount protocol to get the first file handle, which allows them entry into a remote filesystem.

The mount protocol is kept separate from the NFS protocol to make it easy to plug in new access checking and validation methods without changing the NFS server protocol.

Note: the protocol definition implies stateful servers because the server maintains a list of client's mount requests. The mount list information is not critical for the correct functioning of either the client or the server. It is intended for advisory use only, for example, to warn people when a server is going down.

**PROGRAMMING**

**#include <rpcsvc/mount.h>**

The following XDR routines are available in **librpcsvc**:

**xdr\_exportbody**  
**xdr\_exports**  
**xdr\_fhandle**  
**xdr\_fhstatus**  
**xdr\_groups**  
**xdr\_mountbody**  
**xdr\_mountlist**  
**xdr\_path**

**SEE ALSO**

**mount(8), mountd(8C), showmount(8)**

*NFS Protocol Spec*, in *Network Programming*

**NAME**

nlm\_prot – protocol between local and remote network lock managers

**PROTOCOL**

/usr/include/rpcsvc/nlm\_prot.x

**DESCRIPTION**

The network lock manager protocol is used for communication between local and remote lock managers.

**PROGRAMMING**

**#include <rpcsvc/nlm\_prot.h>**

**XDR Routines**

The following XDR routines are available in **librpcsvc**:

**xdr\_nlm\_testargs**

**xdr\_nlm\_testres**

**xdr\_nlm\_lockargs**

**xdr\_nlm\_cancargs**

**xdr\_nlm\_unlockargs**

**xdr\_nlm\_res**

**SEE ALSO**

**lockd(8C)**



**NAME**

pnp – automatic network installation

**PROTOCOL**

`/usr/include/rpcsvc/pnprpc.x`

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

`pnp()` is used during unattended network installation, and routine booting, of Sun386i systems on a Sun386i network. Each network cable (subnetwork or full network) must have at least one `pnpd(8C)` server running on it to support PNP.

**PROGRAMMING**

`#include <rpcsvc/pnprpc.h>`

The following RPC calls are available in version 2 of the PNP protocol:

**NULLPROC**

Finds a PNP daemon on the local network. Used with `clntudp_broadcast()`, often to measure network overhead.

**PNP\_WHOAMI**

Used early in the boot process to acquire network configuration information about a system, or to determine that a system is not known by the network.

**PNP\_ACQUIRE**

Used to acquire a server willing to configure a new system after a `PNP_WHOAMI` request fails. This RPC is typically broadcast; any successful reply may be used.

**PNP\_SETUP**

Requests a network configuration from a PNP daemon that has responded to a previous `PNP_ACQUIRE` RPC.

**PNP\_POLL**

After a `PNP_SETUP` request, if the status is `in_progress`, the procedure is to wait 20 seconds, and issue a `PNP_POLL` request, and then check the status again. Once the status is `success`, the system will be configured for the network. Entries in the yp database may be added or old ones deleted, and file storage may be assigned, according to the architecture and boot type.

If the server misses 5 `PNP_POLL` requests, it will assume that the client system crashed and back out of the procedure. Similarly, if the client system does not receive responses from the server for `PNP_MISSEDPOLLS` consecutive requests, it should assume the server crashed and begin its PNP sequence again.

**SEE ALSO**

`pnpboot(8C)`, `pnpd(8C)`

**NAME**

publickey, getpublickey, getsecretkey – get public or secret key

**SYNOPSIS**

```
#include <rpc/rpc.h>
#include <rpc/key_prot.h>

getpublickey(netname, publickey)
char netname[MAXNETNAMELEN+1];
char publickey[HEXKEYBYTES+1];

getsecretkey(netname, secretkey, passwd)
char netname[MAXNETNAMELEN+1];
char secretkey[HEXKEYBYTES+1];
char *passwd;
```

**DESCRIPTION**

These routines are used to get public and secret keys from the YP database. **getsecretkey()** has an extra argument, *passwd*, which is used to decrypt the encrypted secret key stored in the database. Both routines return 1 if they are successful in finding the key, 0 otherwise. The keys are returned as NULL-terminated, hexadecimal strings. If the password supplied to **getsecretkey()** fails to decrypt the secret key, the routine will return 1 but the *secretkey* argument will be a NULL string.

**SEE ALSO**

**publickey(5)**

*RPC Programmer's Manual in Network Programming*

**NAME**

rex – remote execution protocol

**PROTOCOL**

`/usr/include/rpcsvc/rex.x`

**DESCRIPTION**

This server will execute commands remotely. The working directory and environment of the command can be specified, and the standard input and output of the command can be arbitrarily redirected. An option is provided for interactive I/O for programs that expect to be running on terminals. Note: this service is only provided with the TCP transport.

**PROGRAMMING**

```
#include <sys/ioctl.h>
```

```
#include <rpcsvc/rex.h> /* not compiled with rpgen */
```

The following XDR routines are available in `librpcsvc`:

```
xdr_rex_start()  
xdr_rex_result()  
xdr_rex_tty_mode()  
xdr_rex_tty_size()
```

**SEE ALSO**

`on(1C)`, `rexd(8C)`

**NAME**

**rnusers, rusers** – return information about users on remote machines

**PROTOCOL**

**/usr/include/rpcsvc/rnusers.x**

**DESCRIPTION**

**rnusers()** returns the number of users logged on to *host* (-1 if it cannot determine that number). **rusers()** fills the **utmpidlearr** structure with data about *host*, and returns 0 if successful.

**PROGRAMMING**

```
#include <rpcsvc/rusers.h>
rnusers(host)
char *host
rusers(host, up)
char *host
struct utmpidlearr *up;
```

The following XDR routines are also available:

```
xdr_utmpidle
xdr_utmpidlearr
```

**SEE ALSO**

**rusers(1C)**

**NAME**

**quota** – implement quotas on remote machines

**PROTOCOL**

**/usr/include/rpcsvc/rquota.x**

**DESCRIPTION**

The **rquota()** protocol inquires about quotas on remote machines. It is used in conjunction with NFS, since NFS itself does not implement quotas.

**PROGRAMMING**

**#include <rpcsvc/rquota.h>**

The following XDR routines are available in **librpcsvc**:

**xdr\_getquota\_arg**

**xdr\_getquota\_rslt**

**xdr\_rquota**

**SEE ALSO**

**quota(1), quotactl(2)**

**NAME**

**rstat** – get performance data from remote kernel

**PROTOCOL**

**/usr/include/rpcsvc/rstat.x**

**DESCRIPTION**

The **rstat()** protocol is used to gather statistics from remote kernel. Statistics are available on items such as paging, swapping and cpu utilization.

**PROGRAMMING**

**#include <rpcsvc/rstat.h>**

**havedisk(host)**

**char \*host;**

**rstat(host, statp)**

**char \*host;**

**struct statstime \*statp;**

**havedisk()** returns 1 if *host* has a disk, 0 if it does not, and -1 if this cannot be determined. **rstat()** fills in the **statstime** structure for *host*, and returns 0 if it was successful.

The following XDR routines are available in **librpcsvc**:

**xdr\_statstime**

**xdr\_statsswtch**

**xdr\_stats**

**SEE ALSO**

**perfmeter(1), rup(1C), rstatd(8C)**

**NAME**

*rwall* – write to specified remote machines

**SYNOPSIS**

```
#include <rpcsvc/rwall.h>

rwall(host, msg);
char *host, *msg;
```

**DESCRIPTION**

*host* prints the string *msg* to all its users. It returns 0 if successful.

**RPC INFO****program number:**

WALLPROG

**procs:**

WALLPROC\_WALL

Takes string as argument (*wrapstring*), returns no arguments.  
Executes *wall* on remote host with string.

**versions:**

RSTATVERS\_ORIG

**SEE ALSO**

*rwall*(1C), *rwalld*(8C), *shutdown*(8)

**NAME**

sm\_inter – status monitor protocol

**PROTOCOL**

/usr/include/rpcsvc/sm\_inter.x

**DESCRIPTION**

The status monitor protocol is used for monitoring the status of remote hosts.

**PROGRAMMING**

#include <rpcsvc/sm\_inter.h>

**XDR Routines**

The following XDR routines are available in **librpcsvc**:

xdr\_sm\_name  
xdr\_mon  
xdr\_mon\_id  
xdr\_sm\_stat\_res  
xdr\_sm\_stat

**SEE ALSO**

statd(8C)



**NAME**

spray – scatter data in order to check the network

**PROTOCOL**

`/usr/include/rpcsvc/spray.x`

**DESCRIPTION**

The spray protocol sends packets to a given machine to test the speed and reliability of it.

**PROGRAMMING**

`#include <rpcsvc/spray.h>`

The following XDR routines are available in `librpcsvc`:

`xdr_sprayarr`

`xdr_spraycumul`

**SEE ALSO**

`spray(8C)`, `sprayd(8C)`

**NAME**

**xcrypt, xencrypt, xdecrypt, passwd2des** – hex encryption and utility routines

**SYNOPSIS**

**xencrypt(data, key)**

**char \*data;**

**char \*key;**

**xdecrypt(data, key)**

**char \*data;**

**char \*key;**

**passwd2des(pass, key)**

**char \*pass;**

**char \*key;**

**DESCRIPTION**

The routines **xencrypt** and **xdecrypt** take null-terminated hexadecimal strings as arguments, and encrypt them using the 8-byte *key* as input to the DES algorithm. The input strings must have a length that is a multiple of 16 hex digits (64 bits is the DES block size).

**passwd2des** converts a password, of arbitrary length, into an 8-byte DES key, with odd-parity set in the low bit of each byte. The high-order bit of each input byte is ignored.

These routines are used by the DES authentication subsystem for encrypting and decrypting the secret keys stored in the **publickey** database.

**SEE ALSO**

**des\_crypt(3), publickey(5)**

**NAME**

yp – NIS protocol

**PROTOCOL**

/usr/include/rpcsvc/yp.x

**DESCRIPTION**

The Network Information Service (NIS) is used for the administration of network-wide databases. The service is composed mainly of two programs: **YPBINDPROG** for finding a NIS server and **YPPROG** for accessing the NIS databases.

**PROGRAMMING**

Refer to **ypclnt(3N)** for information on the programmatic interface to NIS servers and databases.

**SEE ALSO**

**ypclnt(3N)**, **yppasswd(3R)**

**NOTES**

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications plc, and may not be used without permission.

**NAME**

yppasswd – update user password in NIS

**PROTOCOL**

/usr/include/rpcsvc/yppasswd.x

**DESCRIPTION**

The `yppasswd()` protocol is used to change a user's password entry in the Network Information Service (NIS) password database.

If *oldpass* is indeed the old user password, this routine replaces the password entry with *newpw*. It returns 0 if successful.

**PROGRAMMING**

```
#include <rpcsvc/yppasswd.h>
```

```
yppasswd(oldpass, newpw)  
char *oldpass  
struct passwd *newpw;
```

**SEE ALSO**

yppasswd(1), yppasswdd(8C)

**NOTES**

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications plc, and may not be used without permission.

**NAME**

intro – introduction to device drivers, protocols, and network interfaces

**DESCRIPTION**

This section describes device drivers, high-speed network interfaces, and protocols available under SunOS. The system provides drivers for a variety of hardware devices, such as disks, magnetic tapes, serial communication lines, mice and frame buffers, as well as virtual devices such as pseudo-terminals and windows. SunOS provides hardware support and a network interface for the 10-Megabit Ethernet, along with interfaces for the IP protocol family and a STREAMS-based Network Interface Tap (NIT) facility.

In addition to describing device drivers that are supported by the 4.3BSD operating system, this section contains subsections that describe:

- SunOS-specific device drivers, under '4S'.
- Protocol families, under '4F'.
- Protocols and raw interfaces, under '4P'.
- STREAMS modules, under '4M'.
- Network interfaces, under '4N'.

**Configuration**

The SunOS kernel can be configured to include or omit many of the device drivers described in this section. The CONFIG section of the manual page gives the line(s) to include in the kernel configuration file for each machine architecture on which a device is supported. If no specific architectures are indicated, the configuration syntax applies to all Sun systems.

The GENERIC kernel is the default configuration for SunOS. It contains all of the optional drivers for a given machine architecture. See `config(8)`, for details on configuring a new SunOS kernel.

The manual page for a device driver may also include a DIAGNOSTICS section, listing error messages that the driver might produce. Normally, these messages are logged to the appropriate system log using the kernel's standard message-buffering mechanism (see `syslogd(8)`); they may also appear on the system console.

**Ioctls**

Various special functions, such as querying or altering the operating characteristics of a device, are performed by supplying appropriate parameters to the `ioctl(2)` system call. These parameters are often referred to as "ioctls." Ioctls for a specific device are presented in the manual page for that device. Ioctls that pertain to a class of devices are listed in a manual page with a name that suggests the class of device, and ending in 'io', such as `mtio(4)` for magnetic tape devices, or `dkio(4S)` for disk controllers. In addition, some ioctls operate directly on higher-level objects such as files, terminals, sockets, and streams:

- Ioctls that operate directly on files, file descriptors, and sockets are described in `filio(4)`. Note: the `fcntl(2V)` system call is the primary method for operating on file descriptors as such, rather than on the underlying files. Also note that the `setsockopt` system call (see `getsockopt(2)`) is the primary method for operating on sockets as such, rather than on the underlying protocol or network interface. Ioctls for a specific network interface are documented in the manual page for that interface.
- Ioctls for terminals, including pseudo-terminals, are described in `termio(4)`. This manual page includes information about both the BSD `termios` structure, as well as the System V `termio` structure.
- Ioctls for STREAMS are described in `streamio(4)`.

**Devices Always Present**

Device drivers present in every kernel include:

- The paging device; see `drum(4)`.
- Drivers for accessing physical, virtual, and I/O space in memory; see `mem(4S)`.
- The data sink; see `null(4)`.

**Terminals and Serial Communications Devices**

Serial communication lines are normally supported by the terminal driver; see **tty(4)**. This driver manages serial lines provided by communications drivers, such as those described in **mti(4S)** and **zs(4S)**. The terminal driver also handles serial lines provided by virtual terminals, such as the Sun console monitor described in **console(4S)**, and true pseudo-terminals, described in **pty(4)**.

**Disk Devices**

Drivers for the following disk controllers provide standard block and raw interfaces under SunOS;

- SCSI controllers, in **sd(4S)**,
- Xylogics 450 and 451 SMD controllers, in **xy(4S)**,
- Xylogics 7053 SMD controllers, in **xd(4S)**.

Ioctls to query or set a disk's geometry and partitioning are described in **dkio(4S)**.

**Magnetic Tape Devices**

Magnetic tape devices supported by SunOS include those described in **ar(4S)**, **tm(4S)**, **st(4S)**, and **xt(4S)**. Ioctls for all tape-device drivers are described in **mtio(4S)**.

**Frame Buffers**

Frame buffer devices include color frame buffers described in the **cg\*(4S)** manual pages, monochrome frame buffers described in the **bw\*(4S)** manual pages, graphics processor interfaces described in the **gp\*(4S)** manual pages, and an indirect device for the console frame buffer described in **fb(4S)**. Ioctls for all frame-buffer devices are described in **fbio(4S)**.

**Miscellaneous Devices**

Miscellaneous devices include the console keyboard described in **kbd(4S)**, the console mouse described in **mouse(4S)**, window devices described in **win(4S)**, and the DES encryption-chip interface described in **des(4S)**.

**Network-Interface Devices**

SunOS supports the 10-Megabit Ethernet as its primary network interface; see **ie(4S)** and **le(4S)** for details. However, a software loopback interface, **lo(4)** is also supported. General properties of these network interfaces are described in **if(4N)**, along with the ioctls that operate on them.

Support for network routing is described in **routing(4N)**.

**Protocols and Protocol Families**

SunOS supports both socket-based and STREAMS-based network communications. The Internet protocol family, described in **inet(4F)**, is the primary protocol family primary supported by SunOS, although the system can support a number of others. The raw interface provides low-level services, such as packet fragmentation and reassembly, routing, addressing, and basic transport for socket-based implementations. Facilities for communicating using an Internet-family protocol are generally accessed by specifying the **AF\_INET** address family when binding a socket; see **socket(2)** for details.

Major protocols in the Internet family include:

- The Internet Protocol (IP) itself, which supports the universal datagram format, as described in **ip(4P)**. This is the default protocol for **SOCK\_RAW** type sockets within the **AF\_INET** domain.
- The Transmission Control Protocol (TCP); see **tcp(4P)**. This is the default protocol for **SOCK\_STREAM** type sockets.
- The User Datagram Protocol (UDP); see **udp(4P)**. This is the default protocol for **SOCK\_DGRAM** type sockets.
- The Address Resolution Protocol (ARP); see **arp(4P)**.
- The Internet Control Message Protocol (ICMP); see **icmp(4P)**.

The Network Interface Tap (NIT) protocol, described in `nit(4P)`, is a STREAMS-based facility for accessing the network at the link level.

## SEE ALSO

`fcntl(2V)`, `getsockopt(2)`, `ioctl(2)`, `socket(2)`, `ar(4S)`, `arp(4P)`, `dkio(4S)`, `drum(4)`, `fb(4S)`, `fbio(4S)`, `filio(4)`, `icmp(4P)`, `if(4N)`, `inet(4F)`, `ip(4P)`, `kbd(4S)`, `le(4S)`, `lo(4)`, `mem(4S)`, `mti(4S)`, `mtio(4)`, `nit(4P)`, `null(4)`, `pty(4)`, `routing(4N)`, `sd(4S)`, `st(4S)`, `streamio(4)`, `tcp(4P)`, `termio(4)`, `tm(4S)`, `tty(4)`, `udp(4P)`, `win(4S)`, `xd(4S)`, `xy(4S)`, `zs(4S)`

## LIST OF DEVICES, INTERFACES AND PROTOCOLS

Name	Appears on Page	Description
<code>alm</code>	<code>mcp(4S)</code>	ALM-2 Asynchronous Line Multiplexer
<code>ar</code>	<code>ar(4S)</code>	Archive 1/4 inch Streaming Tape Drive
<code>arp</code>	<code>arp(4P)</code>	Address Resolution Protocol
<code>atbus</code>	<code>mem(4S)</code>	main memory and bus I/O space
<code>audio</code>	<code>audio(4)</code>	telephone quality audio device
<code>bwtwo</code>	<code>bwtwo(4S)</code>	black and white memory frame buffer
<code>cdromio</code>	<code>cdromio(4S)</code>	CDROM control operations
<code>cgeight</code>	<code>cgeight(4S)</code>	24-bit color memory frame buffer
<code>cgfour</code>	<code>cgfour(4S)</code>	Sun-3 color memory frame buffer
<code>cgnine</code>	<code>cgnine(4S)</code>	24-bit VME color memory frame buffer
<code>cgsix</code>	<code>cgsix(4S)</code>	accelerated 8-bit color frame buffer
<code>cgthree</code>	<code>cgthree(4S)</code>	8-bit color memory frame buffer
<code>cgtwo</code>	<code>cgtwo(4S)</code>	color graphics interface
<code>console</code>	<code>console(4S)</code>	console driver and terminal emulator
<code>db</code>	<code>db(4M)</code>	SunDials STREAMS module
<code>des</code>	<code>des(4S)</code>	DES encryption chip interface
<code>dkio</code>	<code>dkio(4S)</code>	generic disk control operations
<code>drum</code>	<code>drum(4)</code>	paging device
<code>eeprom</code>	<code>mem(4S)</code>	main memory and bus I/O space
<code>fb</code>	<code>fb(4S)</code>	driver for Sun console frame buffer
<code>fbio</code>	<code>fbio(4S)</code>	frame buffer control operations
<code>fd</code>	<code>fd(4S)</code>	Disk driver for Floppy Disk Controllers
<code>filio</code>	<code>filio(4)</code>	ioctls that operate directly on files, file descriptors, and sockets
<code>fpa</code>	<code>fpa(4S)</code>	Sun-3 floating-point accelerator
<code>gpone</code>	<code>gpone(4S)</code>	graphics processor
<code>icmp</code>	<code>icmp(4P)</code>	Internet Control Message Protocol
<code>ie</code>	<code>ie(4S)</code>	Intel 10 Mb/s Ethernet interface
<code>if</code>	<code>if(4N)</code>	general properties of network interfaces
<code>inet</code>	<code>inet(4F)</code>	Internet protocol family
<code>ip</code>	<code>ip(4P)</code>	Internet Protocol
<code>kb</code>	<code>kb(4M)</code>	Sun keyboard STREAMS module
<code>kbd</code>	<code>kbd(4S)</code>	Sun keyboard
<code>kmem</code>	<code>mem(4S)</code>	main memory and bus I/O space
<code>ldterm</code>	<code>ldterm(4M)</code>	standard terminal STREAMS module
<code>le</code>	<code>le(4S)</code>	LANCE 10Mb/s Ethernet interface
<code>lo</code>	<code>lo(4N)</code>	software loopback network interface
<code>lofs</code>	<code>lofs(4S)</code>	loopback virtual file system
<code>mcp</code>	<code>mcp(4S)</code>	MCP Multiprotocol Communications Processor
<code>mem</code>	<code>mem(4S)</code>	main memory and bus I/O space
<code>mouse</code>	<code>mouse(4S)</code>	Sun mouse
<code>ms</code>	<code>ms(4M)</code>	Sun mouse STREAMS module
<code>mti</code>	<code>mti(4S)</code>	Systech MTI-800/1600 multi-terminal interface
<code>mtio</code>	<code>mtio(4)</code>	general magnetic tape interface

<b>NFS</b>	<b>nfs(4P)</b>	network file system
<b>nit</b>	<b>nit(4P)</b>	Network Interface Tap
<b>nit_buf</b>	<b>nit_buf(4M)</b>	STREAMS NIT buffering module
<b>nit_if</b>	<b>nit_if(4M)</b>	STREAMS NIT device interface module
<b>nif_pf</b>	<b>nit_pf(4M)</b>	STREAMS NIT packet filtering module
<b>null</b>	<b>null(4)</b>	data sink
<b>openprom</b>	<b>openprom(4S)</b>	PROM monitor configuration interface
<b>pp</b>	<b>pp(4)</b>	Centronics-compatible parallel printer port
<b>pty</b>	<b>pty(4)</b>	pseudo-terminal driver
<b>rfs</b>	<b>rfs(4)</b>	remote file sharing service
<b>root</b>	<b>root(4S)</b>	pseudo-driver for Sun386i root disk
<b>routing</b>	<b>routing(4N)</b>	system supporting for local network packet routing
<b>sbus</b>	<b>mem(4S)</b>	main memory and bus I/O space
<b>sd</b>	<b>sd(4S)</b>	driver for SCSI disk devices
<b>sockio</b>	<b>sockio(4)</b>	ioctl's that operate directly on sockets
<b>sr</b>	<b>sr(4S)</b>	driver for CDROM SCSI controller
<b>st</b>	<b>st(4S)</b>	driver for SCSI tape devices
<b>streamio</b>	<b>streamio(4)</b>	STREAMS ioctl commands
<b>taac</b>	<b>taac(4S)</b>	Sun applications accelerator
<b>tcp</b>	<b>tcp(4P)</b>	Internet Transmission Control Protocol
<b>tcpfli</b>	<b>tcpfli(4P)</b>	TLI-Conforming TCP Stream-Head
<b>termio</b>	<b>termio(4)</b>	general terminal interface
<b>tfs</b>	<b>tfs(4S)</b>	translucent file service
<b>tm</b>	<b>tm(4S)</b>	Tapemaster 1/2 inch tape controller
<b>tmpfs</b>	<b>tmpfs(4S)</b>	memory based filesystem
<b>ttcompat</b>	<b>ttcompat(4M)</b>	V7 and 4BSD STREAMS compatibility module
<b>tty</b>	<b>tty(4)</b>	controlling terminal interface
<b>udp</b>	<b>udp(4P)</b>	Internet User Datagram Protocol
<b>unix</b>	<b>unix(4F)</b>	UNIX domain protocol family
<b>vd</b>	<b>vd(4)</b>	loadable modules interface
<b>vme16d16</b>	<b>mem(4S)</b>	main memory and bus I/O space
<b>vme16d32</b>	<b>mem(4S)</b>	main memory and bus I/O space
<b>vme24d16</b>	<b>mem(4S)</b>	main memory and bus I/O space
<b>vme24d32</b>	<b>mem(4S)</b>	main memory and bus I/O space
<b>vme32d16</b>	<b>mem(4S)</b>	main memory and bus I/O space
<b>vme32d32</b>	<b>mem(4S)</b>	main memory and bus I/O space
<b>vpc</b>	<b>vpc(4S)</b>	System VPC-2200 Versatec printer/plotter
<b>win</b>	<b>win(4S)</b>	Sun window system
<b>xd</b>	<b>xd(4S)</b>	Disk driver for Xylogics 7053 SMD Disk Controller
<b>xt</b>	<b>xt(4S)</b>	Xylogics 472 1/2 inch tape controller
<b>xy</b>	<b>xy(4S)</b>	Disk driver for Xylogics 450 and 451 SMD Disk Controllers
<b>zero</b>	<b>mem(4S)</b>	main memory and bus I/O space
<b>zero</b>	<b>zero(4S)</b>	source of zeroes
<b>zs</b>	<b>zs(4S)</b>	Zilog 8530 SCC serial communications driver



**NAME**

ar – Archive 1/4 inch Streaming Tape Drive

**AVAILABILITY**

Sun-3 and Sun-4 systems only.

**DESCRIPTION**

The Archive tape controller is a Sun 'QIC-II' interface to an Archive streaming tape drive. It provides a standard tape interface to the device, see `mtio(4)`, with some deficiencies listed under **BUGS** below.

The maximum blocksize for the raw device is limited only by available memory.

**FILES**

`/dev/rar*`

`/dev/nrar*` non-rewinding

**SEE ALSO**

`mtio(4)`

**DIAGNOSTICS**

**ar\*:** would not initialize

**ar\*:** already open

The tape can be opened by only one process at a time

**ar\*:** no such drive

**ar\*:** no cartridge in drive

**ar\*:** cartridge is write protected

**ar:** interrupt from uninitialized controller %x

**ar\*:** many retries, consider retiring this

**ar\*:** %b error at block #

**ar\*:** %b error at block #

**ar:** giving up on Rdy, try

**BUGS**

The tape cannot reverse direction so the `BSF` and `BSR` ioctls are not supported.

The `FSR` ioctl is not supported.

The system will hang if the tape is removed while running.

When using the raw device, the number of bytes in any given transfer must be a multiple of 512 bytes. If it is not, the device driver returns an error.

The driver will only write an EOF mark on close if the last operation was a write, without regard for the mode used when opening the file. This delete empty files on a raw tape copy operation.

## NAME

arp – Address Resolution Protocol

## CONFIG

pseudo-device ether

## SYNOPSIS

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

s = socket(AF_INET, SOCK_DGRAM, 0);
```

## DESCRIPTION

ARP is a protocol used to dynamically map between Internet Protocol (IP) and 10Mb/s Ethernet addresses. It is used by all the 10Mb/s Ethernet interface drivers. It is not specific to the Internet Protocol or to the 10Mb/s Ethernet, but this implementation currently supports only that combination.

ARP caches IP-to-Ethernet address mappings. When an interface requests a mapping for an address not in the cache, ARP queues the message which requires the mapping and broadcasts a message on the associated network requesting the address mapping. If a response is provided, the new mapping is cached and any pending message is transmitted. ARP will queue at most one packet while waiting for a mapping request to be responded to; only the most recently “transmitted” packet is kept.

To facilitate communications with systems which do not use ARP, `ioctl()` requests are provided to enter and delete entries in the IP-to-Ethernet tables.

## USAGE

```
#include <sys/sockio.h>
#include <sys/socket.h>
#include <net/if.h>
#include <net/if_arp.h>
struct arpreq arpreq;
ioctl(s, SIOCSARP, (caddr_t)&arpreq);
ioctl(s, SIOCGARP, (caddr_t)&arpreq);
ioctl(s, SIOCDDARP, (caddr_t)&arpreq);
```

Each `ioctl()` takes the same structure as an argument. `SIOCSARP` sets an ARP entry, `SIOCGARP` gets an ARP entry, and `SIOCDDARP` deletes an ARP entry. These `ioctl()` requests may be applied to any socket descriptor `s`, but only by the super-user. The `arpreq` structure contains:

```
/*
 * ARP ioctl request
 */
struct arpreq {
    struct sockaddr arp_pa;    /* protocol address */
    struct sockaddr arp_ha;    /* hardware address */
    int    arp_flags;         /* flags */
};
/* arp_flags field values */
#define ATF_COM            0x2    /* completed entry (arp_ha valid) */
#define ATF_PERM          0x4    /* permanent entry */
#define ATF_PUBL          0x8    /* publish (respond for other host) */
#define ATF_USETRAILERS   0x10   /* send trailer packets to host */
```

The address family for the `arp_pa` `sockaddr` must be `AF_INET`; for the `arp_ha` `sockaddr` it must be `AF_UNSPEC`. The only flag bits which may be written are `ATF_PERM`, `ATF_PUBL` and `ATF_USETRAILERS`. `ATF_PERM` makes the entry permanent if the `ioctl()` call succeeds. The peculiar nature of the ARP tables may cause the `ioctl()` to fail if more than 6 (permanent) IP addresses hash to the same slot. `ATF_PUBL` specifies that the ARP code should respond to ARP requests for the indicated host

coming from other machines. This allows a host to act as an "ARP server" which may be useful in convincing an ARP-only machine to talk to a non-ARP machine.

ARP is also used to negotiate the use of trailer IP encapsulations; trailers are an alternate encapsulation used to allow efficient packet alignment for large packets despite variable-sized headers. Hosts which wish to receive trailer encapsulations so indicate by sending gratuitous ARP translation replies along with replies to IP requests; they are also sent in reply to IP translation replies. The negotiation is thus fully symmetrical, in that either or both hosts may request trailers. The ATF\_USETRAILERS flag is used to record the receipt of such a reply, and enables the transmission of trailer packets to that host.

ARP watches passively for hosts impersonating the local host (that is, a host which responds to an ARP mapping request for the local host's address).

#### SEE ALSO

**ec(4S), ie(4S), inet(4F), arp(8C), ifconfig(8C)**

Plummer, Dave, "*An Ethernet Address Resolution Protocol -or- Converting Network Protocol Addresses to 48.bit Ethernet Addresses for Transmission on Ethernet Hardware*," RFC 826, Network Information Center, SRI International, Menlo Park, Calif., November 1982. (Sun 800-1059-10)

Leffler, Sam, and Michael Karels, "*Trailer Encapsulations*," RFC 893, Network Information Center, SRI International, Menlo Park, Calif., April 1984.

#### DIAGNOSTICS

**duplicate IP address!! sent from ethernet address: %x:%x:%x:%x:%x:%x.**

ARP has discovered another host on the local network which responds to mapping requests for its own Internet address.

#### BUGS

ARP packets on the Ethernet use only 42 bytes of data, however, the smallest legal Ethernet packet is 60 bytes (not including CRC). Some systems may not enforce the minimum packet size, others will.

**NAME**

audio – telephone quality audio device

**CONFIG**

**device-driver audio**

**AVAILABILITY**

This device is available with SPARCstation 1 systems only.

**DESCRIPTION**

The **audio** device plays and records a single channel of sound using the AM79C30A Digital Subscriber Controller chip. The chip has a built-in analog to digital converter (ADC) and digital to analog converter (DAC) that can drive either the built-in speaker or an external headphone jack, selectable under software control. Digital audio data is sampled at a rate of 8000 samples per second with 12-bit precision, though the data is compressed, using  $\mu$ -law encoding, to 8-bit samples. The resulting audio data quality is equivalent to that of standard telephone service.

The **audio** driver is implemented as a STREAMS device. In order to record audio input, applications **open(2V)** the **/dev/audio** device and read data from it using the **read(2V)** system call. Similarly, sound data is queued to the audio output port by using the **write(2V)** system call.

**Opening the Audio Device**

The audio device is treated as an exclusive resource: only one process may typically open the device at a time. However, two processes may simultaneously access the device if one opens it read-only and the other opens it write-only.

When a process cannot open **/dev/audio** because the requested access mode is busy:

- if the **O\_NDELAY** flag is set in the **open()** *flags* argument, then **open()** returns **-1** immediately, with *errno* set to **EBUSY**.
- if **O\_NDELAY** is not set, then **open()** hangs until the device is available or a signal is delivered to the process, in which case **open()** returns **-1** with *errno* set to **EINTR**.

Since the audio device grants exclusive read or write access to a single process at a time, long-lived audio applications may choose to close the device when they enter an idle state, reopening it when required. The *play.waiting* and *record.waiting* flags in the audio information structure (see below) provide an indication that another process has requested access to the device. This information is advisory only; background audio output processes, for example, may choose to relinquish the audio device whenever another process requests write access.

**Recording Audio Data**

The **read()** system call copies data from the system buffers to the application. Ordinarily, **read()** blocks until the user buffer is filled. The **FIONREAD ioctl** (see **filio(4)**) may be used to determine the amount of data that may be read without blocking. The device may alternatively be set to a non-blocking mode, in which case **read()** completes immediately, but may return fewer bytes than requested. Refer to the **read(2V)** manual page for a complete description of this behavior.

When the audio device is opened with read access, the device driver immediately starts buffering audio input data. Since this consumes system resources, processes that do not record audio data should open the device write-only (**O\_WRONLY**).

The transfer of input data to STREAMS buffers may be paused (or resumed) by using the **AUDIO\_SETINFO ioctl** to set (or clear) the *record.pause* flag in the audio information structure (see below). All unread input data in the STREAMS queue may be discarded by using the **I\_FLUSH STREAMS ioctl** (see **streamio(4)**).

Input data accumulates in STREAMS buffers at a rate of 8000 bytes per second. If the application that consumes the data cannot keep up with this data rate, the STREAMS queue may become full. When this occurs, the *record.error* flag is set in the audio information structure and input sampling ceases until there is room in the input queue for additional data. In such cases, the input data stream contains a discontinuity. For this reason, audio recording applications should open the audio device when they are prepared to begin reading data, rather than at the start of extensive initialization.

**Playing Audio Data**

The `write()` system call copies data from an applications buffer to the STREAMS output queue. Ordinarily, `write()` blocks until the entire user buffer is transferred. The device may alternatively be set to a non-blocking mode, in which case `write()` completes immediately, but may have transferred fewer bytes than requested (see `write(2V)`).

Although `write()` returns when the data is successfully queued, the actual completion of audio output may take considerably longer. The `AUDIO_DRAIN ioctl` may be issued to allow an application to block until all of the queued output data has been played. Alternatively, a process may request asynchronous notification of output completion by writing a zero-length buffer (end-of-file record) to the output stream. When such a buffer has been processed, the `play.eof` flag in the audio information structure (see below) is incremented.

The final `close()` of the file descriptor hangs until audio output has drained. If a signal interrupts the `close()`, or if the process exits without closing the device, any remaining data queued for audio output is flushed and the device is closed immediately.

The conversion of output data may be paused (or resumed) by using the `AUDIO_SETINFO ioctl` to set (or clear) the `play.pause` flag in the audio information structure. Queued output data may be discarded by using the `I_FLUSH STREAMS ioctl`.

Output data is played from the STREAMS buffers at a rate of 8000 bytes per second. If the output queue becomes empty, the `play.error` flag is set in the audio information structure and output ceases until additional data is written.

**Asynchronous I/O**

The `I_SETSIG STREAMS ioctl` may be used to enable asynchronous notification, via the `SIGPOLL` signal, of input and output ready conditions. This, in conjunction with non-blocking `read()` and `write()` requests, is normally sufficient for applications to maintain an audio stream in the background. Alternatively, asynchronous reads and writes may be initiated using the `aioread(3)` functions.

**Audio Data Encoding**

The data samples processed by the audio device are encoded in 8 bits. The high-order bit is a sign bit: 1 represents positive data and 0 represents negative data. The low-order 7 bits represent signal magnitude and are inverted (1's complement). The magnitude is encoded according to a  $\mu$ -law transfer function; such an encoding provides an improved signal-to-noise ratio at low amplitude levels. In order to achieve best results, the audio recording gain should be set so that typical amplitude levels lie within approximately three-fourths of the full dynamic range.

**Audio Control Pseudo-Device**

It is sometimes convenient to have an application, such as a volume control panel, modify certain characteristics of the audio device while it is being used by an unrelated process. The `/dev/audiocntl` minor device is provided for this purpose. Any number of processes may open `/dev/audiocntl` simultaneously. However, `read()` and `write()` system calls are ignored by `/dev/audiocntl`. The `AUDIO_GETINFO` and `AUDIO_SETINFO ioctl` commands may be issued to `/dev/audiocntl` in order to determine the status or alter the behavior of `/dev/audio`.

**Audio Status Change Notification**

Applications that open the audio control pseudo-device may request asynchronous notification of changes in the state of the audio device by setting the `S_MSG` flag in an `I_SETSIG STREAMS ioctl`. Such processes receive a `SIGPOLL` signal when any of the following events occurs:

- An `AUDIO_SETINFO ioctl` has altered the device state.
- An input overflow or output underflow has occurred.
- An end-of-file record (zero-length buffer) has been processed on output.
- An `open()` or `close()` of `/dev/audio` has altered the device state.

### Audio Information Structure

The state of the audio device may be polled or modified using the `AUDIO_GETINFO` and `AUDIO_SETINFO` `ioctl` commands. These commands operate on the `audio_info` structure, defined in `<sun/audioio.h>` as follows:

```

/* Data encoding values, used below in the encoding field */
#define AUDIO_ENCODING_ULAW      (1)    /* u-law encoding */
#define AUDIO_ENCODING_ALAW      (2)    /* A-law encoding */

/* These ranges apply to record, play, and monitor gain values */
#define AUDIO_MIN_GAIN           (0)     /* minimum gain value */
#define AUDIO_MAX_GAIN           (255)   /* maximum gain value */

/* Audio I/O channel status, used below in the audio_info structure */
struct audio_prinfo {
    /* The following values describe the audio data encoding */
    unsigned    sample_rate;    /* samples per second */
    unsigned    channels;       /* number of interleaved channels */
    unsigned    precision;      /* number of bits per sample */
    unsigned    encoding;       /* data encoding method */

    /* The following values control audio device configuration */
    unsigned    gain;           /* gain level */
    unsigned    port;          /* selected I/O port */

    /* The following values describe the current device state */
    unsigned    samples;        /* number of samples converted */
    unsigned    eof;           /* End Of File counter (play only) */
    unsigned char pause;       /* non-zero if paused, zero to resume */
    unsigned char error;       /* non-zero if overflow/underflow */
    unsigned char waiting;     /* non-zero if a process wants access */

    /* The following values are read-only device state flags */
    unsigned char open;        /* non-zero if open access granted */
    unsigned char active;     /* non-zero if I/O active */
};

/* This structure is used in AUDIO_GETINFO and AUDIO_SETINFO ioctl commands */
typedef struct audio_info {
    struct audio_prinfo    record;    /* input status information */
    struct audio_prinfo    play;     /* output status information */
    unsigned               monitor_gain; /* input to output mix */
} audio_info_t;

```

The `play.gain` and `record.gain` fields specify the output and input volume levels. A value of `AUDIO_MAX_GAIN` indicates maximum gain. The device also allows input data to be monitored by mixing audio input onto the output channel. The `monitor.gain` field controls the level of this feedback path. The `play.port` field controls the output path for the audio device. It may be set to either `AUDIO_SPEAKER` or `AUDIO_HEADPHONE` to direct output to the built-in speaker or the headphone jack, respectively.

The `play.pause` and `record.pause` flags may be used to pause and resume the transfer of data between the audio device and the `STREAMS` buffers. The `play.error` and `record.error` flags indicate that data underflow or overflow has occurred. The `play.active` and `record.active` flags indicate that data transfer is currently active in the corresponding direction.

The `play.open` and `record.open` flags indicate that the device is currently open with the corresponding access permission. The `play.waiting` and `record.waiting` flags provide an indication that a process may be waiting to access the device. These flags are set automatically when a process blocks on `open()`, though they may also be set using the `AUDIO_SETINFO` `ioctl` command. They are cleared only when a process relinquishes access by closing the device.

The *play.samples* and *record.samples* fields are initialized, at `open()`, to zero and increment each time a data sample is copied to or from the associated STREAMS queue. Applications that keep track of the number of samples read or written may use these fields to determine exactly how many samples remain in the STREAMS buffers. The *play.eof* field increments whenever a zero-length output buffer is synchronously processed. Applications may use this field to detect the completion of particular segments of audio output.

The *sample\_rate*, *channels*, *precision*, and *encoding* fields report the audio data format in use by the device. For now, these values are read-only; however, future audio device implementations may support more than one data encoding format, in which case applications might be able to modify these fields.

#### Filio and STREAMS IOCTLS

All of the `filio(4)` and `streamio(4)` `ioctl` commands may be issued for the `/dev/audio` device. Because the `/dev/audiocctl` device has its own STREAMS queues, most of these commands neither modify nor report the state of `/dev/audio` if issued for the `/dev/audiocctl` device. The `I_SETSIG` `ioctl` may be issued for `/dev/audiocctl` to enable the notification of audio status changes, as described above.

#### Audio IOCTLS

The audio device additionally supports the following `ioctl` commands:

##### AUDIO\_DRAIN

The argument is ignored. This command suspends the calling process until the output STREAMS queue is empty, or until a signal is delivered to the calling process. It may only be issued for the `/dev/audio` device. An implicit `AUDIO_DRAIN` is performed on the final `close()` of `/dev/audio`.

##### AUDIO\_GETINFO

The argument is a pointer to an `audio_info` structure. This command may be issued for either `/dev/audio` or `/dev/audiocctl`. The current state of the `/dev/audio` device is returned in the structure.

##### AUDIO\_SETINFO

The argument is a pointer to an `audio_info` structure. This command may be issued for either `/dev/audio` or `/dev/audiocctl`. This command configures the audio device according to the structure supplied and overwrites the structure with the new state of the device. [Note: The *play.samples*, *record.samples*, *play.error*, *record.error*, and *play.eof* fields are modified to reflect the state of the device when the `AUDIO_SETINFO` was issued. This allows programs to atomically modify these fields while retrieving the previous value.]

Certain fields in the information structure, such as the *pause* flags, are treated as read-only when `/dev/audio` is not open with the corresponding access permission. Other fields, such as the gain levels and encoding information, may have a restricted set of acceptable values. Applications that attempt to modify such fields should check the returned values to be sure that the corresponding change took effect.

Once set, the following values persist through subsequent `open()` and `close()` calls of the device: *play.gain*, *record.gain*, *monitor.gain*, *play.port*, and *record.port*. All other state is reset when the corresponding I/O stream of `/dev/audio` is closed.

The `audio_info` structure may be initialized through the use of the `AUDIO_INITINFO` macro. This macro sets all fields in the structure to values that are ignored by the `AUDIO_SETINFO` command. For instance, the following code switches the output port from the built-in speaker to the headphone jack without modifying any other audio parameters:

```
audio_info_t    info;

AUDIO_INITINFO(&info);
info.play.port = AUDIO_HEADPHONE;
err = ioctl(audio_fd, AUDIO_SETINFO, &info);
```

This technique is preferred over using a sequence of `AUDIO_GETINFO` followed by `AUDIO_SETINFO`.

**Unsupported Device Control Features**

The AM79C30A chip is capable of performing a number of functions that are not currently supported by the device driver, many of which were designed primarily for telephony applications. For example, the chip can generate ringer tones and has a number of specialized filtering capabilities that are designed to compensate for different types of external speakers and microphones.

Ordinarily, applications do not need to access these capabilities and, further, altering the chip's characteristics may interfere with its normal behavior. However, knowledgeable applications may use the unsupported **AUDIOGETREG** and **AUDIOSETREG** **ioctl** commands to read and write the chip registers directly. The description of this interface may be found in `<sbusedev/audio_79C30.h>`. Note: these commands are supplied for prototyping purposes only and may become obsolete in a future release of the audio driver.

**FILES**

`/dev/audio`  
`/dev/audioctl`  
`/usr/demo/SOUND`

**SEE ALSO**

**ioctl(2)**, **poll(2)**, **read(2V)**, **write(2V)**, **aioread(3)**, **filio(4)**, **streamio(4)**

AMD data sheet for the AM79C30A Digital Subscriber Controller, Publication number 09893.

**BUGS**

Due to a *feature* of the STREAMS implementation, programs that are terminated or exit without closing the **audio** device may hang for a short period while audio output drains. In general, programs that produce audio output should catch the SIGINT signal and flush the output stream before exiting.

The current driver implementation does not support the A-law encoding mode of the AM79C30A chip. Future implementations may permit the **AUDIO\_SETINFO** **ioctl** to modify the *play.encoding* and *record.encoding* fields of the device information structure to enable this mode.

**FUTURE DIRECTIONS**

Workstation audio resources should be managed by a networked audio server, in the same way that the video monitor is manipulated by a window system server. For the time being, we encourage you to write your programs in a modular fashion, isolating the **audio** device-specific functions, so that they may be easily ported to such an environment.



**NAME**

bwtwo – black and white memory frame buffer

**CONFIG — SUN-3, SUN-3x SYSTEMS**

```
device bwtwo0 at obmem 1 csr 0xff000000 priority 4
device bwtwo0 at obmem 2 csr 0x100000 priority 4
device bwtwo0 at obmem 3 csr 0xff000000 priority 4
device bwtwo0 at obmem 4 csr 0xff000000
device bwtwo0 at obmem 7 csr 0xff000000 priority 4
device bwtwo0 at obmem ? csr 0x50300000 priority 4
```

The first synopsis line given above is used to generate a kernel for Sun-3/75, Sun-3/140 or Sun-3/160 systems; the second, for a Sun-3/50 system; the third, for a Sun-3/260 system; the fourth, for a Sun-3/110 system; the fifth, for a Sun-3/60 system; and the sixth for Sun-3/80 and Sun-3/470 systems.

**CONFIG — SUN-4 SYSTEMS**

```
device bwtwo0 at obio 1 csr 0xfd000000 priority 4
device bwtwo0 at obio 2 csr 0xfb300000 priority 4
device bwtwo0 at obio 3 csr 0xfb300000 priority 4
device bwtwo0 at obio 4 csr 0xfb300000 priority 4
```

The first synopsis line given above should be used to generate a kernel for a Sun-4/260 or Sun-4/280 system; the second, for a Sun-4/110 system; the third for a Sun-4/330 system; and the fourth for a Sun-4/460 system.

**CONFIG — SPARCstation 1 SYSTEMS**

```
device-driver bwtwo
```

**CONFIG — Sun386i SYSTEM**

```
device bwtwo0 at obmem ? csr 0xA0200000
```

**DESCRIPTION**

The **bwtwo** interface provides access to Sun monochrome memory frame buffers. It supports the ioctl's described in **fbio(4S)**.

If **flags 0x1** is specified, frame buffer write operations are buffered through regular high-speed RAM. This “copy memory” mode of operation speeds frame buffer accesses, but consumes an extra 128K bytes of memory. Only Sun-3/75, Sun-3/140, and Sun-3/160 systems support copy memory; on other systems a warning message is printed and the flag is ignored.

Reading or writing to the frame buffer is not allowed — you must use the **mmap(2)** system call to map the board into your address space.

**FILES**

```
/dev/bwtwo[0-9]    device files
```

**SEE ALSO**

```
mmap(2), cgfour(4S), fb(4S), fbio(4S)
```

**BUGS**

Use of vertical-retrace interrupts is not supported.

## NAME

cdromio – CDROM control operations

## DESCRIPTION

The Sun CDROM device driver supports a set of `ioctl(2)` commands for audio operations and CDROM specific operations. It also supports the `dkio(4S)` operations — generic disk control operation for all Sun disk drivers. See `dkio(4S)` Basic to these `cdromio ioctl()` requests are the definitions in `<scsi/targets/srdef.h>` or `<sundev/srreg.h>`

```

/*
 * CDROM I/O controls type definitions
 */

/* definition of play audio msf structure */
struct cdrom_msf {
    unsigned char    cdmsf_min0;    /* starting minute */
    unsigned char    cdmsf_sec0;    /* starting second */
    unsigned char    cdmsf_frame0;  /* starting frame */
    unsigned char    cdmsf_min1;    /* ending minute */
    unsigned char    cdmsf_sec1;    /* ending second */
    unsigned char    cdmsf_frame1;  /* ending frame */
};

/* definition of play audio track/index structure */
struct cdrom_ti {
    unsigned char    cdti_trk0;     /* starting track */
    unsigned char    cdti_ind0;     /* starting index */
    unsigned char    cdti_trk1;     /* ending track */
    unsigned char    cdti_ind1;     /* ending index */
};

/* definition of read toc header structure */
struct cdrom_tochdr {
    unsigned char    cdth_trk0;     /* starting track */
    unsigned char    cdth_trk1;     /* ending track */
};

/* definition of read toc entry structure */
struct cdrom_tocentry {
    unsigned char    cdte_track;
    unsigned char    cdte_adr       :4;
    unsigned char    cdte_ctrl      :4;
    unsigned char    cdte_format;
    union {
        struct {
            unsigned char    minute;
            unsigned char    second;
            unsigned char    frame;
        } msf;
        int    lba;
    } cdte_addr;
    unsigned char    cdte_datamode;
};

```

```

/*
 * Bitmask for CDROM data track in the cdte_ctrl field
 * A track is either data or audio.
 */
#define CDROM_DATA_TRACK 0x04

/*
 * CDROM address format definition, for use with struct cdrom_tocentry
 */
#define CDROM_LBA 0x01
#define CDROM_MSF 0x02

/*
 * For CDROMREADTOCENTRY, set the cdte_track to CDROM_LEADOUT to get
 * the information for the leadout track.
 */
#define CDROM_LEADOUT 0xAA

struct cdrom_subchnl {
    unsigned char  cdsc_format;
    unsigned char  cdsc_audiostatus;
    unsigned char  cdsc_adr: 4;
    unsigned char  cdsc_ctrl: 4;
    unsigned char  cdsc_trk;
    unsigned char  cdsc_ind;
    union {
        struct {
            unsigned char  minute;
            unsigned char  second;
            unsigned char  frame;
        } msf;
        int  lba;
    } cdsc_absaddr;
    union {
        struct {
            unsigned char  minute;
            unsigned char  second;
            unsigned char  frame;
        } msf;
        int  lba;
    } cdsc_reladdr;
};

/*
 * Definition for audio status returned from Read Sub-channel
 */
#define CDROM_AUDIO_INVALID 0x00 /* audio status not supported */
#define CDROM_AUDIO_PLAY 0x11 /* audio play operation in progress */
#define CDROM_AUDIO_PAUSED 0x12 /* audio play operation paused */
#define CDROM_AUDIO_COMPLETED 0x13 /* audio play successfully completed */
#define CDROM_AUDIO_ERROR 0x14 /* audio play stopped due to error */
#define CDROM_AUDIO_NO_STATUS 0x15 /* no current audio status to return */

```

```

/* definition of audio volume control structure */
struct cdrom_volctrl {
    unsigned char  cdvc_chnl0;
    unsigned char  cdvc_chnl1;
    unsigned char  cdvc_chnl2;
    unsigned char  cdvc_chnl3;
};

struct cdrom_read {
    int  cunread_lba;
    caddr_t cunread_bufaddr;
    int  cunread_buflen;
};

#define CDROM_MODE1_SIZE    2048
#define CDROM_MODE2_SIZE    2336

/*
 * CDROM I/O control commands
 */
#define CDROMPAUSE  _IO(c, 10) /* Pause Audio Operation */

#define CDROMRESUME _IO(c, 11) /* Resume paused Audio Operation */

#define CDROMPLAYMSF _IOW(c, 12, struct cdrom_msf) /* Play Audio MSF */

#define CDROMPLAYTRKIND _IOW(c, 13, struct cdrom_ti) /* Play Audio Trk/ind */

#define CDROMREADTOCHDR _IOR(c, 103, struct cdrom_tochdr) /* Read TOC hdr */

#define CDROMREADTOCENTRY _IOWR(c, 104, struct cdrom_tocentry) /* Read TOC */

#define CDROMSTOP  _IO(c, 105) /* Stop the cdrom drive */

#define CDROMSTART  _IO(c, 106) /* Start the cdrom drive */

#define CDROMEJECT  _IO(c, 107) /* Ejects the cdrom caddy */

#define CDROMVOLCTRL  _IOW(c, 14, struct cdrom_volctrl) /* volume control */

#define CDROMSUBCHNL _IOWR(c, 108, struct cdrom_subchnl) /* read subchannel */

#define CDROMREADMODE2  _IOW(c, 110, struct cdrom_read) /* mode 2 */

#define CDROMREADMODE1  _IOW(c, 111, struct cdrom_read) /* mode 1 */

```

The CDROMPAUSE ioctl() pauses the current audio play operation and the CDROMRESUME ioctl() resumes the paused audio play operation. The CDROMSTART ioctl() spins up the disc and seeks to the last address requested, while the CDROMSTOP ioctl() spins down the disc and the CDROMEJECT ioctl() ejects the caddy with the disc. All of the above ioctl() calls only take a file descriptor and a command as arguments. They have the form:

```

ioctl(fd, cmd)
    int  fd;
    int  cmd;

```

The rest of the `ioctl()` calls have the form:

```
ioctl(fd, cmd, ptr)
    int    fd;
    int    cmd;
    char   *ptr;
```

where `ptr` is a pointer to a struct or an integer.

The `CDROMPLAYMSF ioctl()` command requests the drive to output the audio signals starting at the specified starting address and continue the audio play until the specified ending address is detected. The address is in MSF (minute, second, frame) format. The third argument of the function call is a pointer to the type `struct cdrom_msf`.

The `CDROMPLAYTRKIND ioctl()` command is similar to `CDROMPLAYMSF`. The starting and ending address is in track/index format. The third argument of the function call is a pointer to the type `struct cdrom_ti`.

The `CDROMREADTOCHDR ioctl()` command returns the header of the TOC (table of contents). The header consists of the starting tracking number and the ending track number of the disc. These two numbers are returned through a pointer of `struct cdrom_tochdr`. While the disc can start at any number, all tracks between the first and last tracks are in contiguous ascending order. A related `ioctl()` command is `CDROMREADTOCENTRY`. This command returns the information of a specified track. The third argument of the function call is a pointer to the type `struct cdrom_tocentry`. The caller need to supply the track number and the address format. This command will return a 4-bit `adr` field, a 4-bit `ctrl` field, the starting address in MSF format or LBA format, and the data mode if the track is a data track. The `ctrl` field specifies whether the track is data or audio. To get information for the lead-out area, supply the `ioctl()` command with the track field set to `CDROM_LEADOUT (0xAA)`.

The `CDROMVOLCTRL ioctl()` command controls the audio output level. The SCSI command allows the control of up to 4 channels. The current implementation of the supported CDROM drive only uses channel 0 and channel 1. The valid values of volume control are between 0x00 and 0xFF, with a value of 0xFF indicating maximum volume. The third argument of the function call is a pointer to `struct cdrom_volctrl` which contains the output volume values.

The `CDROMSUBCHNL ioctl()` command reads the Q sub-channel data of the current block. The sub-channel data includes track number, index number, absolute CDROM address, track relative CDROM address, control data and audio status. All information is returned through a pointer to `struct cdrom_subchnl`. The caller needs to supply the address format for the returned address.

The `CDROMREADMODE2` and `CDROMREADMODE1 ioctl()` commands are only available on SPARCstation 1 systems.

Finally, on SPARCstation 1 systems only, the driver supports the user SCSI command interface. By issuing the `ioctl()` command, `USCSICMD`, The caller can supply any SCSI-2 commands that the CDROM drive supports. The caller has to provide all the parameters in the SCSI command block, as well as other information such as the user buffer address and buffer length. See the definitions in `<scsi/impl/uscsi.h>`. The `ioctl()` call has the form:

```
ioctl(fd, cmd, ptr)
    int    fd;
    int    cmd;
    char   *ptr;
```

where *ptr* is a pointer to the type:

```
struct uscsi_scmd {
    caddr_t uscsi_cdb;
    int     uscsi_cdblen;
    caddr_t uscsi_bufaddr;
    int     uscsi_buflen;
    unsigned char uscsi_status;
    int     uscsi_flags;
};
```

**uscsi\_cdb** is a pointer to the SCSI command block. Group 0 **cdb**'s are 6 bytes long while the other groups are 10 bytes or 12 bytes. **uscsi\_cdblen** is the length of the **cdb**. **uscsi\_bufaddr** is the pointer to the user buffer for parameter passing or data input/output. **buflen** is the length of the user buffer. **uscsi\_flags** are the execution flags for SCSI input/output. The possible flags are **USCSI\_SILENT**, **USCSI\_DIAGNOSE**, **USCSI\_ISOLATE**, **USCSI\_READ**, and **USCSI\_WRITE**.

#### FILES

```
/usr/include/scsi/targets/srdef.h
/usr/include/scsi/impl/uscsi.h
/usr/include/sundev/srreg.h
```

#### SEE ALSO

**ioctl(2)**, **dkio(4S)**, **sr(4S)**

#### BUGS

The interface to this device is preliminary and subject to change in future releases. You are encouraged to write your programs in a modular fashion so that you can easily incorporate future changes.

**NAME**

**cgeight** – 24-bit color memory frame buffer

**CONFIG — SUN-3 AND SUN-4 SYSTEMS**

**device cgeight0 at obmem 7 csr 0xff300000 priority 4**  
**device cgeight0 at obio 4 csr 0xfb300000 priority 4**

The first synopsis line should be used to generate a kernel for the Sun-3/60; the second synopsis for a Sun-4/110 or Sun-4/150 system.

**CONFIG — SUN-3x SYSTEM**

**device cgeight0 at obio ? csr 0x50300000 priority 4**

**DESCRIPTION**

The **cgeight** is a 24-bit color memory frame buffer with a monochrome overlay plane and an overlay enable plane implemented optionally on the Sun-4/110, Sun-4/150, Sun-3/60, Sun-3/470 and Sun-3/80 system models. It provides the standard frame buffer interface as defined in **fbio(4S)**.

In addition to the ioctls described under **fbio(4S)**, the **cgeight** interface responds to two **cgeight**-specific colormap ioctls, **FBIOPUTCMAP** and **FBIOGETCMAP**. **FBIOPUTCMAP** returns no information other than success/failure using the ioctl return value. **FBIOGETCMAP** returns its information in the arrays pointed to by the red, green, and blue members of its **fbcmmap** structure argument; **fbcmmap** is defined in `<sun/fbio.h>` as:

```

struct fbcmmap {
    int          index;          /* first element (0 origin) */
    int          count;         /* number of elements */
    unsigned char *red;         /* red color map elements */
    unsigned char *green;      /* green color map elements */
    unsigned char *blue;       /* blue color map elements */
};

```

The driver uses color board vertical-retrace interrupts to load the colormap.

The systems have an overlay plane colormap, which is accessed by encoding the plane group into the index value with the `PIX_GROUP` macro (see `<pixrect/pr_planegroups.h>`).

When using the `mmap` system call to map in the **cgeight** frame buffer. The device looks like:

<b>DACBASE: 0x200000</b>	<b>-&gt; Brooktree Ramdac</b>	<b>16 bytes</b>
<b>0x202000</b>	<b>-&gt; P4 Register</b>	<b>4 bytes</b>
<b>OVLBASE: 0x210000</b>	<b>-&gt; Overlay Plane</b>	<b>1152x900x1</b>
<b>0x230000</b>	<b>-&gt; Overlay Enable Planea</b>	<b>1152x900x1</b>
<b>0x250000</b>	<b>-&gt; 24-bit Frame Buffera</b>	<b>1152x900x32</b>

**FILES**

`/dev/cgeight0`  
`<sun/fbio.h>`  
`<pixrect/pr_planegroups.h>`

**SEE ALSO**

**mmap(2)**, **fbio(4S)**

**NAME**

cgfour – Sun-3 color memory frame buffer

**CONFIG — SUN-3 SYSTEMS**

device cgfour0 at obmem 4 csr 0xff000000 priority 4  
 device cgfour0 at obmem 7 csr 0xff300000 priority 4

The first synopsis line given should be used to generate a kernel for the Sun-3/110 system; and the second, for a Sun-3/60 system.

**CONFIG — SUN-3x SYSTEMS**

device cgfour0 at obmem ? csr 0x50300000 priority 4

**CONFIG — SUN-4 SYSTEMS**

device cgfour0 at obio 2 csr 0xfb300000 priority 4  
 device cgfour0 at obio 3 csr 0xfb300000 priority 4  
 device cgfour0 at obio 4 csr 0xfb300000 priority 4

The first synopsis line given should be used to generate a kernel for the Sun-4/110 system; the second, for a Sun-4/330 system; and the third for a Sun-4/460 system.

**DESCRIPTION**

The **cgfour** is a color memory frame buffer with a monochrome overlay plane and an overlay enable plane implemented on the Sun-3/110 system and some Sun-3/60 system models. It provides the standard frame buffer interface as defined in **fbio(4S)**.

In addition to the ioctls described under **fbio(4S)**, the **cgfour** interface responds to two **cgfour**-specific colormap ioctls, **FBIOPUTCMAP** and **FBIOGETCMAP**. **FBIOPUTCMAP** returns no information other than success/failure using the ioctl return value. **FBIOGETCMAP** returns its information in the arrays pointed to by the red, green, and blue members of its **fbcmmap** structure argument; **fbcmmap** is defined in `<sun/fbio.h>` as:

```

struct fbcmmap {
    int          index;          /* first element (0 origin) */
    int          count;         /* number of elements */
    unsigned char *red;         /* red color map elements */
    unsigned char *green;      /* green color map elements */
    unsigned char *blue;       /* blue color map elements */
};

```

The driver uses color board vertical-retrace interrupts to load the colormap.

The Sun-3/60 system has an overlay plane colormap, which is accessed by encoding the plane group into the index value with the `PIX_GROUP` macro (see `<pixrect/pr_plane groups.h>`).

**FILES**

`/dev/cgfour0`

**SEE ALSO**

`mmap(2)`, `fbio(4S)`



**NAME**

**cgnine** – 24-bit VME color memory frame buffer

**CONFIGURATION**

**device cgnine0 at vme32d32 ? csr 0x08000000 priority 4 vector cgnineintr 0xaa**

**DESCRIPTION**

**cgnine** is a 24-bit double-buffered VME-based color frame buffer. It provides the standard frame buffer interface defined in **fbio(4S)**, and can be paired with the GP2 graphics accelerator board using **gpconfig(8)**.

**cgnine** has two bits of overlay planes, each of which is a 1-bit deep frame buffer that overlays the 24-bit plane group. When either bit of the two overlay planes is non-zero, the pixel shows the color of the overlay plane. If both bits are zero, the color frame buffer underneath is visible.

The 24-bit frame buffer pixel is organized as one longword (32 bits) per pixel. The pixel format is defined in `<pixrect/pixrect.h>` as follows:

```

union fbunit {
    unsigned int    packed; /* whole-sale deal */
    struct {
        unsigned int    A:8; /* unused, for now */
        unsigned int    B:8; /* blue channel */
        unsigned int    G:8; /* green channel */
        unsigned int    R:8; /* red channel */
    }
        channel; /* access per channel */
};

```

When the board is in double-buffer mode, the low 4 bits of each channel are ignored when written to, which yields 12-bit double-buffering.

The higher bit of the overlay planes ranges from offset 0 to 128K (0x20000) bytes. The lower bit ranges from 128K to 256K bytes. The 4MB (0x400000) of the 24-bit deep pixels begins at 256K. The addresses of the control registers start at the next page after the 24-bit deep pixels.

**FILES**

<b>/dev/cgnine0</b>	device special file
<b>/dev/gpone0a</b>	<b>cgnine</b> bound with GP2
<b>/dev/fb</b>	default frame buffer

**SEE ALSO**

**mmap(2)**, **fbio(4S)**, **gpone(4S)** **gpconfig(8)**

**NAME**

**cgsix** – accelerated 8-bit color frame buffer

**CONFIG — SUN-3, SUN-3x, SUN-4 SYSTEMS**

**device cgsix0 at obmem ? csr 0xff000000 priority 4**

**device cgsix0 at obmem ? csr 0x50000000 priority 4**

**device cgsix0 at obio ? csr 0xfb000000 priority 4**

The first synopsis line given should be used for Sun-3/60 systems, the second for Sun-3x systems, and the third for Sun-4 systems.

**CONFIG — SPARCstation 1 SYSTEMS**

**device-driver cgsix**

**DESCRIPTION**

The **cgsix** is a low-end graphics accelerator designed to enhance vector and polygon drawing performance. It has an 8-bit color frame buffer and provides the standard frame buffer interface as defined in **fbio(4S)**.

The **cgsix** has registers and memory that may be mapped with **mmap(2)**, using the offsets defined in **<sundev/cg6reg.h>**.

**FILES**

**/dev/cgsix0**

**SEE ALSO**

**mmap(2)**, **fbio(4S)**

**NAME**

cgthree – 8-bit color memory frame buffer

**CONFIG — SPARCstation 1 SYSTEMS**

device-driver cgthree

**CONFIG — Sun386i SYSTEM**

device cgthree0 at obmem ? csr 0xA0400000

**AVAILABILITY**

SPARCstation 1 and Sun386i systems only.

**DESCRIPTION**

cgthree is a color memory frame buffer. It provides the standard frame buffer interface as defined in fbio(4S).

**FILES**

/dev/cgthree[0-9]

**SEE ALSO**

mmap(2), fbio(4S)

**NAME**

`cgtwo` – color graphics interface

**CONFIG — SUN-3, SUN-3x, SUN-4 SYSTEMS**

`cgtwo0` at `vme24d16` ? `csr 0x400000` `priority 4` `vector cgtwointr 0xa8`

**DESCRIPTION**

The `cgtwo` interface provides access to the color graphics controller board, which is normally supplied with a 19" 66 Hz non-interlaced color monitor. It provides the standard frame buffer interface as defined in `fbio(4S)`.

The hardware consumes 4 megabytes of VME bus address space. The board starts at standard address `0x400000`. The board must be configured for interrupt level 4.

**FILES**

`/dev/cgtwo[0-9]`

**SEE ALSO**

`mmap(2)`, `fbio(4S)`

**NAME**

**clone** – open any minor device on a STREAMS driver

**DESCRIPTION**

**clone** is a STREAMS software driver that finds and opens an unused minor device on another STREAMS driver. The minor device passed to **clone** during the open operation is interpreted as the major device number of another STREAMS driver for which an unused minor device is to be obtained. Each such open results in a separate stream to a previously unused minor device.

The **clone** driver supports only an **open(2V)** function. This open function performs all of the necessary work so that subsequent system calls (including **close(2V)**) require no further involvement of the **clone** driver.

**ERRORS**

**clone** generates an ENXIO error, without opening the device, if the minor device number provided does not correspond to a valid major device, or if the driver indicated is not a STREAMS driver.

**WARNINGS**

Multiple opens of the same minor device are not supported through the **clone** interface. Executing **stat(2V)** on the file system node for a cloned device yields a different result than does executing **fstat** using a file descriptor obtained from opening that node.

**SEE ALSO**

**close(2V)**, **open(2V)**, **stat(2V)**

**NAME**

console – console driver and terminal emulator for the Sun workstation

**CONFIG**

None; included in standard system.

**SYNOPSIS**

```
#include <fcntl.h>
#include <sys/termios.h>
open("/dev/console", mode);
```

**DESCRIPTION**

**console** is an indirect driver for the Sun console terminal. On a Sun workstation, this driver refers to the workstation console driver, which implements a standard UNIX system terminal. On a Sun server without a keyboard or a frame buffer, this driver refers to the CPU serial port driver (**zs(4S)**); a terminal is normally connected to this port.

The workstation console does not support any of the **termio(4)** device control functions specified by flags in the **c\_cflag** word of the **termios** structure or by the **IGNBRK**, **IGNPAR**, **PARMRK**, or **INPCK** flags in the **c\_iflag** word of the **termios** structure, as these functions apply only to asynchronous serial ports. All other **termio(4)** functions must be performed by STREAMS modules pushed atop the driver; when a slave device is opened, the **ldterm(4M)** and **ttcompat(4M)** STREAMS modules are automatically pushed on top of the stream, providing the standard **termio(4)** interface.

The workstation console driver calls the PROM resident monitor to output data to the console frame buffer. Keystrokes from the CPU serial port to which the keyboard is connected are routed through the keyboard STREAMS module (**kb(4M)**) and treated as input.

When the Sun window system **win(4S)** is active, console input is directed through the window system rather than being treated as input by the workstation console driver.

**IOCTLS**

An ioctl **TIOCCONS** can be applied to pseudo-terminals (**pty(4)**) to route output that would normally appear on the console to the pseudo-terminal instead. Thus, the window system does a **TIOCCONS** on a pseudo-terminal so that the system will route console output to the window to which that pseudo-terminal is connected, rather than routing output through the PROM monitor to the screen, since routing output through the PROM monitor destroys the integrity of the screen. Note: when you use **TIOCCONS** in this way, the console *input* is routed from the pseudo-terminal as well.

If a **TIOCCONS** is performed on **/dev/console**, or the pseudo-terminal to which console output is being routed is closed, output to the console will again be routed to the workstation console driver.

**ANSI STANDARD TERMINAL EMULATION**

The Sun Workstation's PROM monitor provides routines that emulates a standard ANSI X3.64 terminal.

Note: the VT100 also follows the ANSI X3.64 standard but both the Sun and the VT100 have nonstandard extensions to the ANSI X3.64 standard. The Sun terminal emulator and the VT100 are *not* compatible in any true sense.

The Sun console displays 34 lines of 80 ASCII characters per line, with scrolling, (x, y) cursor addressability, and a number of other control functions.

The Sun console displays a non-blinking block cursor which marks the current line and character position on the screen. ASCII characters between 0x20 (space) and 0x7E (tilde) inclusive are printing characters — when one is written to the Sun console (and is not part of an escape sequence), it is displayed at the current cursor position and the cursor moves one position to the right on the current line. If the cursor is already at the right edge of the screen, it moves to the first character position on the next line. If the cursor is already at the right edge of the screen on the bottom line, the Line-feed function is performed (see CTRL-J below), which scrolls the screen up by one or more lines or wraps around, before moving the cursor to the first character position on the next line.

### Control Sequence Syntax

The Sun console defines a number of control sequences which may occur in its input. When such a sequence is written to the Sun console, it is not displayed on the screen, but effects some control function as described below, for example, moves the cursor or sets a display mode.

Some of the control sequences consist of a single character. The notation

`CTRL-X`

for some character *X*, represents a control character.

Other ANSI control sequences are of the form

`ESC [ paramschar`

Spaces are included only for readability; these characters must occur in the given sequence without the intervening spaces.

`ESC` represents the ASCII escape character (ESC, CTRL-[, 0x1B).

`[` The next character is a left square bracket '[' (0x5B).

*params* are a sequence of zero or more decimal numbers made up of digits between 0 and 9, separated by semicolons.

*char* represents a function character, which is different for each control sequence.

Some examples of syntactically valid escape sequences are (again, ESC represent the single ASCII character 'Escape'):

<code>ESC[m</code>	<i>select graphic rendition with default parameter</i>
<code>ESC[7m</code>	<i>select graphic rendition with reverse image</i>
<code>ESC[33;54H</code>	<i>set cursor position</i>
<code>ESC[123;456;0;;3;B</code>	<i>move cursor down</i>

Syntactically valid ANSI escape sequences which are not currently interpreted by the Sun console are ignored. Control characters which are not currently interpreted by the Sun console are also ignored.

Each control function requires a specified number of parameters, as noted below. If fewer parameters are supplied, the remaining parameters default to 1, except as noted in the descriptions below.

If more than the required number of parameters is supplied, only the last *n* are used, where *n* is the number required by that particular command character. Also, parameters which are omitted or set to zero are reset to the default value of 1 (except as noted below).

Consider, for example, the command character *M* which requires one parameter. `ESC[;M` and `ESC[0M` and `ESC[M` and `ESC[23;15;32;1M` are all equivalent to `ESC[1M` and provide a parameter value of 1. Note: `ESC[;5M` (interpreted as 'ESC[5M') is *not* equivalent to `ESC[5;M` (interpreted as 'ESC[5;1M') which is ultimately interpreted as 'ESC[1M').

In the syntax descriptions below, parameters are represented as '#' or '#1;#2'.

### ANSI Control Functions

The following paragraphs specify the ANSI control functions implemented by the Sun console. Each description gives:

- the control sequence syntax
- the hex equivalent of control characters where applicable
- the control function name and ANSI or Sun abbreviation (if any).
- description of parameters required, if any
- description of the control function
- for functions which set a mode, the initial setting of the mode. The initial settings can be restored with the SUNRESET escape sequence.

**Control Character Functions****CTRL-G (0x7) Bell (BEL)**

The Sun Workstation Model 100 and 100U is not equipped with an audible bell. It 'rings the bell' by flashing the entire screen. The window system flashes the window.

**CTRL-H (0x8) Backspace (BS)**

The cursor moves one position to the left on the current line. If it is already at the left edge of the screen, nothing happens.

**CTRL-I (0x9) Tab (TAB)**

The cursor moves right on the current line to the next tab stop. The tab stops are fixed at every multiple of 8 columns. If the cursor is already at the right edge of the screen, nothing happens; otherwise the cursor moves right a minimum of one and a maximum of eight character positions.

**CTRL-J (0xA) Line-feed (LF)**

The cursor moves down one line, remaining at the same character position on the line. If the cursor is already at the bottom line, the screen either scrolls up or "wraps around" depending on the setting of an internal variable *S* (initially 1) which can be changed by the ESC[*r* control sequence. If *S* is greater than zero, the entire screen (including the cursor) is scrolled up by *S* lines before executing the line-feed. The top *S* lines scroll off the screen and are lost. *S* new blank lines scroll onto the bottom of the screen. After scrolling, the line-feed is executed by moving the cursor down one line.

If *S* is zero, 'wrap-around' mode is entered. 'ESC [ 1 r' exits back to scroll mode. If a line-feed occurs on the bottom line in wrap mode, the cursor goes to the same character position in the top line of the screen. When any line-feed occurs, the line that the cursor moves to is cleared. This means that no scrolling occurs. Wrap-around mode is not implemented in the window system.

The screen scrolls as fast as possible depending on how much data is backed up waiting to be printed. Whenever a scroll must take place and the console is in normal scroll mode ('ESC [ 1 r'), it scans the rest of the data awaiting printing to see how many line-feeds occur in it. This scan stops when any control character from the set {VT, FF, SO, SI, DLE, DC1, DC2, DC3, DC4, NAK, SYN, ETB, CAN, EM, SUB, ESC, FS, GS, RS, US} is found. At that point, the screen is scrolled by *N* lines ( $N \geq 1$ ) and processing continues. The scanned text is still processed normally to fill in the newly created lines. This results in much faster scrolling with scrolling as long as no escape codes or other control characters are intermixed with the text.

See also the discussion of the 'Set scrolling' (ESC[*r*) control function below.

**CTRL-K (0xB) Reverse Line-feed**

The cursor moves up one line, remaining at the same character position on the line. If the cursor is already at the top line, nothing happens.

**CTRL-L (0xC) Form-feed (FF)**

The cursor is positioned to the Home position (upper-left corner) and the entire screen is cleared.

**CTRL-M (0xD) Return (CR)**

The cursor moves to the leftmost character position on the current line.

**Escape Sequence Functions****CTRL-[ (0x1B) Escape (ESC)**

This is the escape character. Escape initiates a multi-character control sequence.

**ESC[#@ Insert Character (ICH)**

Takes one parameter, # (default 1). Inserts # spaces at the current cursor position. The tail of the current line starting at the current cursor position inclusive is shifted to the right by # character positions to make room for the spaces. The rightmost # character positions shift off the line and are lost. The position of the cursor is unchanged.



- ESC[#A**            **Cursor Up (CUU)**  
 Takes one parameter, # (default 1). Moves the cursor up # lines. If the cursor is fewer than # lines from the top of the screen, moves the cursor to the topmost line on the screen. The character position of the cursor on the line is unchanged.
- ESC[#B**            **Cursor Down (CUD)**  
 Takes one parameter, # (default 1). Moves the cursor down # lines. If the cursor is fewer than # lines from the bottom of the screen, move the cursor to the last line on the screen. The character position of the cursor on the line is unchanged.
- ESC[#C**            **Cursor Forward (CUF)**  
 Takes one parameter, # (default 1). Moves the cursor to the right by # character positions on the current line. If the cursor is fewer than # positions from the right edge of the screen, moves the cursor to the rightmost position on the current line.
- ESC[#D**            **Cursor Backward (CUB)**  
 Takes one parameter, # (default 1). Moves the cursor to the left by # character positions on the current line. If the cursor is fewer than # positions from the left edge of the screen, moves the cursor to the leftmost position on the current line.
- ESC[#E**            **Cursor Next Line (CNL)**  
 Takes one parameter, # (default 1). Positions the cursor at the leftmost character position on the #-th line below the current line. If the current line is less than # lines from the bottom of the screen, positions the cursor at the leftmost character position on the bottom line.
- ESC[#1;#2f**        **Horizontal And Vertical Position (HVP)**  
 or  
**ESC[#1;#2H**        **Cursor Position (CUP)**  
 Takes two parameters, #1 and #2 (default 1, 1). Moves the cursor to the #2-th character position on the #1-th line. Character positions are numbered from 1 at the left edge of the screen; line positions are numbered from 1 at the top of the screen. Hence, if both parameters are omitted, the default action moves the cursor to the home position (upper left corner). If only one parameter is supplied, the cursor moves to column 1 of the specified line.
- ESC[J**             **Erase in Display (ED)**  
 Takes no parameters. Erases from the current cursor position inclusive to the end of the screen. In other words, erases from the current cursor position inclusive to the end of the current line and all lines below the current line. The cursor position is unchanged.
- ESC[K**             **Erase in Line (EL)**  
 Takes no parameters. Erases from the current cursor position inclusive to the end of the current line. The cursor position is unchanged.
- ESC[#L**            **Insert Line (IL)**  
 Takes one parameter, # (default 1). Makes room for # new lines starting at the current line by scrolling down by # lines the portion of the screen from the current line inclusive to the bottom. The # new lines at the cursor are filled with spaces; the bottom # lines shift off the bottom of the screen and are lost. The position of the cursor on the screen is unchanged.
- ESC[#M**            **Delete Line (DL)**  
 Takes one parameter, # (default 1). Deletes # lines beginning with the current line. The portion of the screen from the current line inclusive to the bottom is scrolled upward by # lines. The # new lines scrolling onto the bottom of the screen are filled with spaces; the # old lines beginning at the cursor line are deleted. The position of the cursor on the screen is unchanged.
- ESC[#P**            **Delete Character (DCH)**  
 Takes one parameter, # (default 1). Deletes # characters starting with the current cursor position. Shifts to the left by # character positions the tail of the current line from the current cursor position inclusive to the end of the line. Blanks are shifted into the rightmost # character positions. The position of the cursor on the screen is unchanged.

- ESC[#m**                    **Select Graphic Rendition (SGR)**  
 Takes one parameter, # (default 0). Note: unlike most escape sequences, the parameter defaults to zero if omitted. Invokes the graphic rendition specified by the parameter. All following printing characters in the data stream are rendered according to the parameter until the next occurrence of this escape sequence in the data stream. Currently only two graphic renditions are defined:
- 0 Normal rendition.
  - 7 Negative (reverse) image.
- Negative image displays characters as white-on-black if the screen mode is currently black-on-white, and vice-versa. Any non-zero value of # is currently equivalent to 7 and selects the negative image rendition.
- ESC[p**                    **Black On White (SUNBOW)**  
 Takes no parameters. Sets the screen mode to black-on-white. If the screen mode is already black-on-white, has no effect. In this mode spaces display as solid white, other characters as black-on-white. The cursor is a solid black block. Characters displayed in negative image rendition (see 'Select Graphic Rendition' above) is white-on-black in this mode. This is the initial setting of the screen mode on reset.
- ESC[q**                    **White On Black (SUNWOB)**  
 Takes no parameters. Sets the screen mode to white-on-black. If the screen mode is already white-on-black, has no effect. In this mode spaces display as solid black, other characters as white-on-black. The cursor is a solid white block. Characters displayed in negative image rendition (see 'Select Graphic Rendition' above) is black-on-white in this mode. The initial setting of the screen mode on reset is the alternative mode, black on white.
- ESC[#r**                    **Set scrolling (SUNSCRL)**  
 Takes one parameter, # (default 0). Sets to # an internal register which determines how many lines the screen scrolls up when a line-feed function is performed with the cursor on the bottom line. A parameter of 2 or 3 introduces a small amount of "jump" when a scroll occurs. A parameter of 34 clears the screen rather than scrolling. The initial setting is 1 on reset.
- A parameter of zero initiates "wrap mode" instead of scrolling. In wrap mode, if a linefeed occurs on the bottom line, the cursor goes to the same character position in the top line of the screen. When any linefeed occurs, the line that the cursor moves to is cleared. This means that no scrolling ever occurs. 'ESC [ 1 r' exits back to scroll mode.
- For more information, see the description of the Line-feed (CTRL-J) control function above.
- ESC[s**                    **Reset terminal emulator (SUNRESET)**  
 Takes no parameters. Resets all modes to default, restores current font from PROM. Screen and cursor position are

#### 4014 TERMINAL EMULATION

The PROM monitor for Sun models 100U and 150U provides the Sun Workstation with the capability to emulate a subset of the Tektronix 4014 terminal. This feature does not exist in other Sun PROMs and will be removed from models 100U and 150U in future Sun releases. `tektool(1)` provides Tektronix 4014 terminal emulation and should be used instead of relying on the capabilities of the PROM monitor.

#### FILES

`/dev/console`

#### SEE ALSO

`tektool(1)`, `kb(4M)`, `ldterm(4M)`, `pty(4)`, `termio(4)`, `ttcompat(4M)`, `win(4S)`, `zs(4S)`

ANSI Standard X3.64, "Additional Controls for Use with ASCII", Secretariat: CBEMA, 1828 L St., N.W., Washington, D.C. 20036.

**BUGS**

**TIOCCONS should be restricted to the owner of */dev/console*.**

**NAME**

db – SunDials STREAMS module

**CONFIG**

pseudo-device db

**SYNOPSIS**

```
#include <sys/stream.h>
#include <sundev/vuid_event.h>
#include <sundev/dbio.h>
#include <sys/time.h>
#include <sys/ioctl.h>
open("/dev/dialbox", O_RDWR);
ioctl(fd, I_PUSH, "db");
```

**DESCRIPTION**

The **db** STREAMS module processes the byte streams generated by the SunDials dial box. The dial box generates a stream of bytes that encode the identity of the dials and the amount by which they are turned.

Each dial sample in the byte stream consists of three bytes. The first byte identifies which dial was turned and the next two bytes return the delta in signed binary format. When bound to an application using the window system, *Virtual User Input Device* events are generated. An event from a dial is constrained to lie between 0x80 and 0x87.

A stream with **db** pushed into it can emit *firm\_events* as specified by the protocol of a VUID. **db** understands the **VIDSFORMAT** and **VIDGFORMAT** ioctls (see reference below), as defined in `/usr/include/sundev/dbio.h` and `/usr/include/sundev/vuid_event.h`. All other `ioctl()` requests are passed downstream. **db** sets the parameters of a serial port when it is opened. No `termios(4)` `ioctl()` requests should be performed on a **db** STREAMS module, as **db** expects the device parameters to remain as it set them.

**IOCTLS**

**VIDSFORMAT**

**VIDGFORMAT**

These are standard *Virtual User Input Device* ioctls. See *SunView System Programmer's Guide* for a description of their operation.

**FILES**

```
/usr/include/sundev/dbio.h
/usr/include/sundev/vuid_event.h
/usr/include/sys/ioctl.h
/usr/include/sys/stream.h
/usr/include/sys/time.h
```

**SEE ALSO**

`termios(4)`, `dialtest(6)`, `dbconfig(8)`

*SunView System Programmer's Guide*,  
*SunDials Programmers Guide*

**BUGS**

**VIDSADDR** and **VIDGADDR** are not supported.

**WARNING**

The SunDials dial box must be used with a serial port.

**NAME**

des – DES encryption chip interface

**CONFIG — SUN-3 SYSTEM**

device des0 at obio ? csr 0x1c0000

**CONFIG — SUN-3x SYSTEM**

device des0 at obio ? csr 0x66002000

**CONFIG — SUN-4 SYSTEM**

device des0 at obio ? csr 0xfe000000

**SYNOPSIS**

**#include <sys/des.h>**

**DESCRIPTION**

The **des** driver provides a high level interface to the AmZ8068 Data Ciphering Processor, a hardware implementation of the NBS Data Encryption Standard.

The high level interface provided by this driver is hardware independent and could be shared by future drivers in other systems.

The interface allows access to two modes of the DES algorithm: Electronic Code Book (ECB) and Cipher Block Chaining (CBC). All access to the DES driver is through **ioctl(2)** calls rather than through reads and writes; all encryption is done in-place in the user's buffers.

**IOCTLS**

The **ioctls** provided are:

**DESIOCBLOCK**

This call encrypts/decrypts an entire buffer of data, whose address and length are passed in the '**struct desparams**' addressed by the argument. The length must be a multiple of 8 bytes.

**DESIOCQUICK**

This call encrypts/decrypts a small amount of data quickly. The data is limited to **DES\_QUICKLEN** bytes, and must be a multiple of 8 bytes. Rather than being addresses, the data is passed directly in the '**struct desparams**' argument.

**FILES**

**/dev/des**

**SEE ALSO**

**des(1)**, **des\_crypt(3)**

*Federal Information Processing Standards Publication 46*

*AmZ8068 DCP Product Description, Advanced Micro Devices*

## NAME

dkio – generic disk control operations

## DESCRIPTION

All Sun disk drivers support a set of `ioctl(2)` requests for disk formatting and labeling operations. Basic to these `ioctl()` requests are the definitions in `/usr/include/sun/dkio.h`:

```

/*
 * Structures and definitions for disk I/O control commands
 */
/* Controller and disk identification */
struct dk_info {
    int     dki_ctlr;        /* controller address */
    short   dki_unit;       /* unit (slave) address */
    short   dki_ctype;      /* controller type */
    short   dki_flags;      /* flags */
};
/* controller types */
#define DKC_UNKNOWN      0
#define DKC_DSD5215     5
#define DKC_XY450       6
#define DKC_ACB4000     7
#define DKC_MD21        8
#define DKC_XD7053     11
#define DKC_CSS         12
#define DKC_NEC765     13    /* floppy on Sun386i */
#define DKC_INTEL82072  14
/* flags */
#define DKI_BAD144    0x01 /* use DEC std 144 bad sector fwding */
#define DKI_MAPTRK   0x02 /* controller does track mapping */
#define DKI_FMTTRK   0x04 /* formats only full track at a time */
#define DKI_FMTVOL   0x08 /* formats only full volume at a time */
/* Definition of a disk's geometry */
struct dk_geom {
    unsigned short   dkg_ncyl;    /* # of data cylinders */
    unsigned short   dkg_acyl;    /* # of alternate cylinders */
    unsigned short   dkg_bcyl;    /* cyl offset (for fixed head area) */
    unsigned short   dkg_nhead;   /* # of heads */
    unsigned short   dkg_bhead;   /* head offset (for Larks, etc.) */
    unsigned short   dkg_nsect;   /* # of sectors per track */
    unsigned short   dkg_intrlv;  /* interleave factor */
    unsigned short   dkg_gap1;    /* gap 1 size */
    unsigned short   dkg_gap2;    /* gap 2 size */
    unsigned short   dkg_apc;     /* alternates per cyl (SCSI only) */
    unsigned short   dkg_extra[9]; /* for compatible expansion */
};
/* Partition map (part of dk_label) */
struct dk_map {
    long   dkl_cylno;    /* starting cylinder */
    long   dkl_nblk;    /* number of blocks */
};

```

```

/* Floppy characteristics */
struct fdk_char {
    u_char medium;      /* medium type (scsi floppy only) */
    int transfer_rate; /* transfer rate */
    int ncyl;           /* number of cylinders */
    int nhead;          /* number of heads */
    int sec_size;       /* sector size */
    int secptrack;      /* sectors per track */
    int steps;          /* number of steps per */
};
/* Used by FDKGETCHANGE, returned state of the sense disk change bit. */
#define FDKGC_HISTORY 0x01 /* disk has changed since last call */
#define FDKGC_CURRENT 0x02 /* current state of disk change */
/* disk I/O control commands */
#define DKIOCINFO      _IOR(d, 8, struct dk_info) /* Get info */
#define DKIOCGGGEOM    _IOR(d, 2, struct dk_geom) /* Get geometry */
#define DKIOCSGGEOM    _IOW(d, 3, struct dk_geom) /* Set geometry */
#define DKIOCGPART     _IOR(d, 4, struct dk_map) /* Get partition info */
#define DKIOCSPART     _IOW(d, 5, struct dk_map) /* Set partition info */
#define DKIOCWCHK      _IOWR(d, 115, int) /* Toggle write check */
/* floppy I/O control commands */
#define FDKIOGCHAR     _IOR(d, 114, struct fdk_char) /* Get floppy characteristics */
#define FDKEJECT       _IO(d, 112) /* Eject floppy */
#define FDKGETCHANGE   _IOR(d, 111, int) /* Get disk change status */

```

The **DKIOCINFO** ioctl returns a **dk\_info** structure which tells the type of the controller and attributes about how bad-block processing is done on the controller. The **DKIOCGPART** and **DKIOCSPART** get and set the controller's current notion of the partition table for the disk (without changing the partition table on the disk itself), while the **DKIOCGGGEOM** and **DKIOCSGGEOM** ioctls do similar things for the per-drive geometry information. The **DKIOCWCHK** enables or disables a disk's write check capabilities. The **FDKIOGCHAR** ioctl returns an **fdk\_char** structure which gives the characteristics of the floppy diskette. The **FDKEJECT** ioctl ejects the floppy diskette. The **FDKGETCHANGE** returns the status of the diskette changed signal from the floppy interface.

#### FILES

/usr/include/sun/dkio.h

#### SEE ALSO

**fd(4S)**, **ip(4P)**, **sd(4S)**, **xd(4S)**, **xy(4S)**, **dkctl(8)**

**NAME**

drum – paging device

**CONFIG**

None; included with standard system.

**SYNOPSIS**

```
#include <fcntl.h>
```

```
open("/dev/drum", mode);
```

**DESCRIPTION**

This file refers to the paging device in use by the system. This may actually be a subdevice of one of the disk drivers, but in a system with paging interleaved across multiple disk drives it provides an indirect driver for the multiple drives.

**FILES**

/dev/drum

**BUGS**

Reads from the drum are not allowed across the interleaving boundaries. Since these only occur every .5Mbytes or so, and since the system never allocates blocks across the boundary, this is usually not a problem.



**NAME**

**fb** – driver for Sun console frame buffer

**CONFIG**

None; included in standard system.

**DESCRIPTION**

The **fb** driver provides indirect access to a Sun frame buffer. It is an indirect driver for the Sun workstation console's frame buffer. At boot time, the workstation's frame buffer device is determined from information from the PROM monitor and set to be the one that **fb** will indirect to. The device driver for the console's frame buffer must be configured into the kernel so that this indirect driver can access it.

The idea behind this driver is that user programs can open a known device, query its characteristics and access it in a device dependent way, depending on the type. **fb** redirects **open(2V)**, **close(2V)**, **ioctl(2)**, and **mmap(2)** calls to the real frame buffer. All Sun frame buffers support the same general interface; see **fbio(4S)**.

**FILES**

**/dev/fb**

**SEE ALSO**

**close(2V)**, **ioctl(2)**, **mmap(2)**, **open(2V)**, **fbio(4S)**

**NAME**

**fbio** – frame buffer control operations

**DESCRIPTION**

All Sun frame buffers support the same general interface that is defined by `<sun/fbio.h>`. Each responds to an **FBIOGTYPE** `ioctl(2)` request which returns information in a **fbtype** structure.

Each device has an **FBTYPE** which is used by higher-level software to determine how to perform graphics functions. Each device is used by opening it, doing an **FBIOGTYPE** `ioctl()` to see which frame buffer type is present, and thereby selecting the appropriate device-management routines.

Full-fledged frame buffers (that is, those that run SunView1) implement an **FBIOPIXRECT** `ioctl()` request, which returns a **pixrect**. This call is made only from inside the kernel. The returned **pixrect** is used by `win(4S)` for cursor tracking and colormap loading.

**FBIOSVIDEO** and **FBIOGVIDEO** are general-purpose `ioctl()` requests for controlling possible video features of frame buffers. These `ioctl()` requests either set or return the value of a flags integer. At this point, only the **FBVIDEO\_ON** option is available, controlled by **FBIOSVIDEO**. **FBIOGVIDEO** returns the current video state.

The **FBIOSATTR** and **FBIOGATTR** `ioctl()` requests allow access to special features of newer frame buffers. They use the **fbattr** and **fbgattr** structures.

Some color frame buffers support the **FBIOPUTCMAP** and **FBIOGETCMAP** `ioctl()` requests, which provide access to the colormap. They use the **fbcmmap** structure.

**SEE ALSO**

`ioctl(2)`, `mmap(2)`, `bw*(4S)`, `cg*(4S)`, `gp*(4S)`, `fb(4S)`, `win(4S)`

**BUGS**

The **FBIOSATTR** and **FBIOGATTR** `ioctl()` requests are only supported by frame buffers which emulate older frame buffer types. For example, `cgfour(4S)` frame buffers emulate `bwtwo(4S)` frame buffers. If a frame buffer is emulating another frame buffer, **FBIOGTYPE** returns the emulated type. To get the real type, use **FBIOGATTR**.

**NAME**

fd – disk driver for Floppy Disk Controllers

**CONFIG — Sun386i SYSTEMS**

controller fdc0 at atmem ? csr 0x1000 dmachan 2 irq 6 priority 2  
disk fd0 at fdc0 drive 0 flags 0

**CONFIG — SUN-3/80 SYSTEMS**

controller fdc0 at obio ? csr 0x6e000000 priority 6 vector fdintr 0x5c  
disk fd0 at fdc0 drive 0 flags 0

**CONFIG — SPARCstation 1 SYSTEMS**

device-driver fd

**AVAILABILITY**

Sun386i, Sun-3/80, and SPARCstation 1 systems only.

**DESCRIPTION**

The **fd** driver provides an interface to floppy disks using the Intel 82072 disk controller on Sun386i, Sun-3/80 and SPARCstation 1 systems.

The minor device number in files that use the floppy interface encodes the unit number as well as the partition. The bits of the minor device number are defined as **rrruuppp** where **r**=reserved, **u**=unit, and **p**=partition. The unit number selects a particular floppy drive for the controller. The partition number picks one of eight partitions [**a-h**].

When the floppy is first opened the driver looks for a label in logical block 0 of the diskette. If a label is found, the geometry and partition information from the label will be used on each access thereafter. The driver first assumes high density characteristics when it tries to read the label. If the read fails it will try the read again using low density characteristics. If both attempts to read the label fail, the open will fail. Use the **FNDELAY** flag when opening an unformatted diskette as a signal to the driver that it should not attempt to access the diskette. If block 0 is read successfully, but a label is not found, the open will fail for the block interface. Using the raw interface, the open will succeed even if the diskette is unlabeled. Default geometry and partitioning are assumed if the diskette is unlabeled.

The default partitions are:

- a**        -> 0, N-1
- b**        -> N-1, N
- c**        -> 0, N

where N is the number of cylinders on the diskette.

The **fd** driver supports both block and raw interfaces. The block files access the disk using the system's normal buffering mechanism and may be read and written without regard to physical disk records. There is also a "raw" interface that provides for direct transmission between the disk and the user's read or write buffer. A single **read(2V)** or **write(2V)** call usually results in one I/O operation; therefore raw I/O is considerably more efficient when many words are transmitted. The names of the raw files conventionally begin with an extra 'r'.

**FILES — Sun386i SYSTEMS**

1.44 MB Floppy Disk Drives:

<b>/dev/fd0a</b>	block file
<b>/dev/fd0c</b>	block file
<b>/dev/rfd0a</b>	raw file
<b>/dev/rfd0c</b>	raw file

## 720 K Floppy Disk Drives:

<b>/dev/fd10a</b>	block file
<b>/dev/fd10c</b>	block file
<b>/dev/rfd10a</b>	raw file
<b>/dev/rfd10c</b>	raw file

## FILES — SUN-3/80 and SPARCstation 1 SYSTEMS

Note: the **fd** driver on Sun-3/80 and SPARCstation 1 systems auto-senses the density of the floppy.

<b>/dev/fd0[a-c]</b>	block file
<b>/dev/fd0</b>	block file (same as <b>/dev/fd0c</b> )
<b>/dev/rfd0[a-c]</b>	raw file
<b>/dev/rfd0</b>	raw file (same as <b>/dev/rfd0c</b> )

## SEE ALSO

**read(2V)**, **write(2V)**, **dkio(4S)**

## DIAGNOSTICS — Sun386i SYSTEMS

**fd drv %d, trk %d: %s**

A command such as **read** or **write** encountered a format-related error condition. The value of *%s* is derived from the error number given by the controller, indicating the nature of the error. The track number is relative to the beginning of the partition involved.

**fd drv %d, blk %d: %s**

A command such as **read** or **write** encountered an error condition related to I/O. The value of *%s* is derived from the error number returned by the controller and indicates the nature of the error. The block number is relative to the start of the partition involved.

**fd controller: %s**

An error occurred in the controller. The value of *%s* is derived from the status returned by the controller and specifies the error encountered.

**fd(%d):%s please insert**

I/O was attempted while the floppy drive door was not latched. The value of *%s* indicates which disk was expected to be in the drive.

## DIAGNOSTICS — SUN-3/80 and SPARCstation 1 SYSTEMS

**fd %d: %s failed (%x %x %x)**

The command, *%s*, failed after several retries on drive *%d*. The three hex values in parenthesis are the contents of status register 0, status register 1, and status register 2 of the Intel 82072 Floppy Disk Controller on completion of the command as documented in the data sheet for that part. This error message is usually followed by one of the following, interpreting the bits of the status register:

**fd %d: not writable**  
**fd %d: crc error**  
**fd %d: overrun/underrun**  
**fd %d: bad format**  
**fd %d: timeout**

## NOTES

Floppy diskettes have 18 sectors per track, and can cross a track (though not a cylinder) boundary without losing data, so when using **dd(1)** to or from a diskette, you should specify **bs=18k** or multiples thereof.

**NAME**

filio – ioctls that operate directly on files, file descriptors, and sockets

**SYNOPSIS**

```
#include <sys/filio.h>
```

**DESCRIPTION**

The IOCTL's listed in this manual page apply directly to files, file descriptors, and sockets, independent of any underlying device or protocol.

Note: the `fcntl(2V)` system call is the primary method for operating on file descriptors as such, rather than on the underlying files.

**IOCTLS for File Descriptors**

**FIOCLEX** The argument is ignored. Set the close-on-exec flag for the file descriptor passed to `ioctl`. This flag is also manipulated by the `F_SETFD` command of `fcntl(2V)`.

**FIONCLEX** The argument is ignored. Clear the close-on-exec flag for the file descriptor passed to `ioctl`.

**IOCTLs for Files**

**FIONREAD** The argument is a pointer to a `long`. Set the value of that `long` to the number of immediately readable characters from whatever the descriptor passed to `ioctl` refers to. This works for files, pipes, sockets, and terminals.

**FIONBIO** The argument is a pointer to an `int`. Set or clear non-blocking I/O. If the value of that `int` is a 1 (one) the descriptor is set for non-blocking I/O. If the value of that `int` is a 0 (zero) the descriptor is cleared for non-blocking I/O.

**FIOASYNC** The argument is a pointer to an `int`. Set or clear asynchronous I/O. If the value of that `int` is a 1 (one) the descriptor is set for asynchronous I/O. If the value of that `int` is a 0 (zero) the descriptor is cleared for asynchronous I/O.

**FIOSETOWN** The argument is a pointer to an `int`. Set the process-group ID that will subsequently receive `SIGIO` or `SIGURG` signals for the object referred to by the descriptor passed to `ioctl` to the value of that `int`.

**FIOGETOWN** The argument is a pointer to an `int`. Set the value of that `int` to the process-group ID that is receiving `SIGIO` or `SIGURG` signals for the object referred to by the descriptor passed to `ioctl`.

**SEE ALSO**

`ioctl(2)`, `fcntl(2V)`, `getsockopt(2)`, `sockio(4)`

**NAME**

fpa – Sun-3/Sun-3x floating-point accelerator

**CONFIG — SUN-3/SUN-3X SYSTEMS**

device fpa0 at virtual ? csr 0xe000000

**SYNOPSIS**

```
#include <sundev/fpareg.h>
open("/dev/fpa", flags);
```

**DESCRIPTION**

FPA and FPA+ are compatible floating point accelerators available on certain Sun-3 and Sun-3x systems. They provide hardware contexts for simultaneous use by up to 32 processes. The same fpa device driver manages either FPA or FPA+ hardware.

Processes access the device using `open(2V)` and `close(2V)` system calls, and the FPA is automatically mapped into the process' address space by SunOS. This is normally provided transparently at compile time by a compiler option, such as the `-ffpa` option to `cc(1V)`.

The valid `ioctl(2)` system calls are used only by diagnostics and by system administration programs, such as `fpa_download(8)`.

**IOCTLS**

<code>FPA_ACCESS_OFF</code>	Clear <code>FPA_ACCESS_BIT</code> in FPA state register to disable access to constants RAM using FPA load pointer.
<code>FPA_ACCESS_ON</code>	Set <code>FPA_ACCESS_BIT</code> in FPA state register to enable access to constants RAM using FPA load pointer.
<code>FPA_FAIL</code>	Disable the FPA.
<code>FPA_GET_DATAREGS</code>	Return the contents of 8 FPA registers.
<code>FPA_INIT_DONE</code>	Called when downloading is complete. Allows multiple users to access the FPA.
<code>FPA_LOAD_OFF</code>	Set <code>FPA_LOAD_BIT</code> in FPA state register to disable access to microstore or map RAM via FPA load pointer.
<code>FPA_LOAD_ON</code>	Set <code>FPA_LOAD_BIT</code> in FPA state register to enable access to microstore or map RAM using FPA load pointer.

The following two `ioctl()` requests are for diagnostic use only. `fpa` must be compiled with `FPA_DIAGNOSTICS_ONLY` defined to enable these two calls.

<code>FPA_WRITE_STATE</code>	Overwrite the FPA state register.
<code>FPA_WRITE_HCP</code>	Write to the hard clear pipe register.

**ERRORS**

The following error messages are returned by `open` system calls only.

<code>EBUSY</code>	All 32 FPA contexts are being used.
<code>EEXIST</code>	The current process has already opened <code>/dev/fpa</code> .
<code>EIO</code>	Downloading has not completed, so only 1 root process can have the FPA open at a time.
<code>ENETDOWN</code>	FPA is disabled.
<code>ENOENT</code>	68881 chip does not exist.
<code>ENXIO</code>	FPA board does not exist.

The following error messages are returned by `ioctl` system calls only.

<code>EINVAL</code>	Invalid <code>ioctl</code> . This may occur if diagnostic only <code>ioctls</code> , <code>FPA_WRITE_STATE</code> or <code>FPA_WRITE_HCP</code> , are used with a driver which didn't compile in those calls.
---------------------	---

**EPERM** All ioctl calls except for `FPA_GET_DATAREGS` require root execution level.  
**EPIPE** The FPA pipe is not clear.

**FILES**

`/dev/fpa` device file for both FPA and FPA+.

**SEE ALSO**

`cc(1V)`, `close(2V)`, `ioctl(2)`, `open(2V)` `fpa_download(8)`, `fparel(8)`, `fpaversion(8)`

**DIAGNOSTICS**

If hardware problems are detected then all processes with `/dev/fpa` open are killed, and future opens of `/dev/fpa` are disabled.

**NAME**

gpone – graphics processor

**CONFIG — SUN-3, SUN-3x, SUN-4 SYSTEMS**

```
device gpone0 at vme24d16 ? csr 0x210000    # GP or GP+
device gpone0 at vme24d32 ? csr 0x240000    # GP2
```

**DESCRIPTION**

The **gpone** interface provides access to the optional Graphics Processor Board (GP).

The hardware consumes 64 kilobytes of VME bus address space. The GP board starts at standard address 0x210000 and must be configured for interrupt level 4.

**IOCTLS**

The graphics processor responds to a number of ioctl calls as described here. One of the calls uses a **gp1fbinfo** structure that looks like this:

```
struct gp1fbinfo {
    int          fb_vmeaddr;    /* physical color board address */
    int          fb_hwwidth;   /* fb board width */
    int          fb_hwheight;  /* fb board height */
    int          addrdelta;    /* phys addr diff between fb and gp */
    caddr_t     fb_ropaddr;    /* cg2 va thru kernelmap */
    int          fbunit;       /* fb unit to use for a,b,c,d */
};
```

The ioctl call looks like this:

```
ioctl(file, request, argp)
int file, request;
```

**argp** is defined differently for each GP ioctl request and is specified in the descriptions below.

The following ioctl commands provide for transferring data between the graphics processor and color boards and processes.

**GP1IO\_PUT\_INFO**

Passes information about the frame buffer into driver. **argp** points to a **struct gp1fbinfo** which is passed to the driver.

**GP1IO\_GET\_STATIC\_BLOCK**

Hands out a static block from the GP. **argp** points to an **int** which is returned from the driver.

**GP1IO\_FREE\_STATIC\_BLOCK**

Frees a static block from the GP. **argp** points to an **int** which is passed to the driver.

**GP1IO\_GET\_GBUFFER\_STATE**

Checks to see if there is a buffer present on the GP. **argp** points to an **int** which is returned from the driver.

**GP1IO\_CHK\_GP**

Restarts the GP if necessary. **argp** points to an **int** which is passed to the driver.

**GP1IO\_GET\_RESTART\_COUNT**

Returns the number of restarts of a GP since power on. Needed to differentiate SIGXCPU calls in user processes. **argp** points to an **int** which is returned from the driver.

**GP1IO\_REDIRECT\_DEVFB**

Configures **/dev/fb** to talk to a graphics processor device. **argp** points to an **int** which is passed to the driver.

**GP1IO\_GET\_REQDEV**

Returns the requested minor device. **argp** points to a **dev\_t** which is returned from the driver.

**GP1IO\_GET\_TRUMINORDEV**



Returns the true minor device. **argp** points to a **char** which is returned from the driver.

The graphics processor driver also responds to the **FBIOGTYPE**, **ioctl** which a program can use to inquire as to the characteristics of the display device, the **FBIOGINFO**, **ioctl** for passing generic information, and the **FBIOGPIXRECT** **ioctl** so that SunWindows can run on it. See **fbio(4S)**.

**FILES**

**/dev/fb**

**/dev/gpone[0-3][abcd]**

**SEE ALSO**

**fbio(4S)**, **mmap(2)**, **gpconfig(8)**

*SunCGI Reference Manual*

**DIAGNOSTICS**

**The Graphics Processor has been restarted. You may see display garbage as a result.**

**NAME**

icmp – Internet Control Message Protocol

**SYNOPSIS**

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

s = socket(AF_INET, SOCK_RAW, proto);
```

**DESCRIPTION**

ICMP is the error and control message protocol used by the Internet protocol family. It is used by the kernel to handle and report errors in protocol processing. It may also be accessed through a “raw socket” for network monitoring and diagnostic functions. The protocol number for ICMP, used in the *proto* parameter to the socket call, can be obtained from **getprotobyname** (see **getprotoent(3N)**). ICMP sockets are connectionless, and are normally used with the *sendto* and *recvfrom* calls, though the **connect(2)** call may also be used to fix the destination for future packets (in which case the **read(2V)** or **recv(2)** and **write(2V)** or **send(2)** system calls may be used).

Outgoing packets automatically have an Internet Protocol (IP) header prepended to them. Incoming packets are provided to the holder of a raw socket with the IP header and options intact.

ICMP is an unreliable datagram protocol layered above IP. It is used internally by the protocol code for various purposes including routing, fault isolation, and congestion control. Receipt of an ICMP “redirect” message will add a new entry in the routing table, or modify an existing one. ICMP messages are routinely sent by the protocol code. Received ICMP messages may be reflected back to users of higher-level protocols such as TCP or UDP as error returns from system calls. A copy of all ICMP message received by the system is provided using the ICMP raw socket.

**ERRORS**

A socket operation may fail with one of the following errors returned:

EISCONN	when trying to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected;
ENOTCONN	when trying to send a datagram, but no destination address is specified, and the socket hasn't been connected;
ENOBUFS	when the system runs out of memory for an internal data structure;
EADDRNOTAVAIL	when an attempt is made to create a socket with a network address for which no network interface exists.

**SEE ALSO**

**connect(2)**, **read(2V)**, **recv(2)**, **send(2)**, **write(2V)**, **getprotoent(3N)**, **inet(4F)**, **ip(4P)**, **routing(4N)**

Postel, Jon, *Internet Control Message Protocol — DARPA Internet Program Protocol Specification*, RFC 792, Network Information Center, SRI International, Menlo Park, Calif., September 1981. (Sun 800-1064-01)

**BUGS**

Replies to ICMP “echo” messages which are source routed are not sent back using inverted source routes, but rather go back through the normal routing mechanisms.

**NAME**

ie – Intel 10 Mb/s Ethernet interface

**CONFIG — SUN-4 SYSTEM**

device ie0 at obio ? csr 0x6000000 priority 3  
 device ie1 at vme24d16 ? csr 0xe88000 priority 3 vector ieintr 0x75  
 device ie2 at vme24d16 ? csr 0x31ff02 priority 3 vector ieintr 0x76  
 device ie3 at vme24d16 ? csr 0x35ff02 priority 3 vector ieintr 0x77

**CONFIG — SUN-3x SYSTEM**

device ie0 at obio ? csr 0x65000000 priority 3  
 device ie1 at vme24d16 ? csr 0xe88000 priority 3 vector ieintr 0x75  
 device ie2 at vme24d32 ? csr 0x31ff02 priority 3 vector ieintr 0x76  
 device ie3 at vme24d32 ? csr 0x35ff02 priority 3 vector ieintr 0x77

**CONFIG — SUN-3 SYSTEM**

device ie0 at obio ? csr 0xc0000 priority 3  
 device ie1 at vme24d16 ? csr 0xe88000 priority 3 vector ieintr 0x75  
 device ie2 at vme24d32 ? csr 0x31ff02 priority 3 vector ieintr 0x76  
 device ie3 at vme24d32 ? csr 0x35ff02 priority 3 vector ieintr 0x77

**CONFIG — SUN-3E SYSTEM**

device ie0 at vme24d16 ? csr 0x31ff02 priority 3 vector ieintr 0x74

**CONFIG — SUN386i SYSTEM**

device ie0 at obmem ? csr 0xD0000000 irq 21 priority 3

**DESCRIPTION**

The **ie** interface provides access to a 10 Mb/s Ethernet network through a controller using the Intel 82586 LAN Coprocessor chip. For a general description of network interfaces see **if(4N)**.

**ie0** specifies a CPU-board-resident interface, except on a Sun-3E where **ie0** is the Sun-3/E Ethernet expansion board. **ie1** specifies a Multibus Intel Ethernet interface for use with a VME adapter. **ie2** and **ie3** specify SunNet Ethernet/VME Controllers, also known as a Sun-3/E Ethernet expansion boards.

**SEE ALSO**

**if(4N)**, **ie(4S)**

**DIAGNOSTICS**

There are too many driver messages to list them all individually here. Some of the more common messages and their meanings follow.

**ie%d: Ethernet jammed**

Network activity has become so intense that sixteen successive transmission attempts failed, and the 82586 gave up on the current packet. Another possible cause of this message is a noise source somewhere in the network, such as a loose transceiver connection.

**ie%d: no carrier**

The 82586 has lost input to its carrier detect pin while trying to transmit a packet, causing the packet to be dropped. Possible causes include an open circuit somewhere in the network and noise on the carrier detect line from the transceiver.

**ie%d: lost interrupt: resetting**

The driver and 82586 chip have lost synchronization with each other. The driver recovers by resetting itself and the chip.

**ie%d: iebark reset**

The 82586 failed to complete a watchdog timeout command in the allotted time. The driver recovers by resetting itself and the chip.

**ie%d: WARNING: requeuing**

The driver has run out of resources while getting a packet ready to transmit. The packet is put back on the output queue for retransmission after more resources become available.

**ie%d: panic: scb overwritten**

The driver has discovered that memory that should remain unchanged after initialization has become corrupted. This error usually is a symptom of a bad 82586 chip.

**ie%d: giant packet**

Provided that all stations on the Ethernet are operating according to the Ethernet specification, this error "should never happen," since the driver allocates its receive buffers to be large enough to hold packets of the largest permitted size. The most likely cause of this message is that some other station on the net is transmitting packets whose lengths exceed the maximum permitted for Ethernet.

**NAME**

if – general properties of network interfaces

**DESCRIPTION**

Each network interface in a system corresponds to a path through which messages may be sent and received. A network interface usually has a hardware device associated with it, though certain interfaces such as the loopback interface, `lo(4)`, do not.

At boot time, each interface with underlying hardware support makes itself known to the system during the autoconfiguration process. Once the interface has acquired its address, it is expected to install a routing table entry so that messages can be routed through it. Most interfaces require some part of their address specified with an `SIOCSIFADDR` IOCTL before they will allow traffic to flow through them. On interfaces where the network-link layer address mapping is static, only the network number is taken from the ioctl; the remainder is found in a hardware specific manner. On interfaces which provide dynamic network-link layer address mapping facilities (for example, 10Mb/s Ethernets using `arp(4P)`), the entire address specified in the ioctl is used.

The following ioctl calls may be used to manipulate network interfaces. Unless specified otherwise, the request takes an `ifreq` structure as its parameter. This structure has the form

```

struct ifreq {
    char    ifr_name[16];          /* name of interface (e.g. "ec0") */
    union {
        struct sockaddr ifru_addr;
        struct sockaddr ifru_dstaddr;
        short   ifru_flags;
    } ifr_ifru;
#define ifr_addr          ifr_ifru.ifru_addr      /* address */
#define ifr_dstaddr      ifr_ifru.ifru_dstaddr   /* other end of p-to-p link */
#define ifr_flags        ifr_ifru.ifru_flags    /* flags */
};

```

<b>SIOCSIFADDR</b>	Set interface address. Following the address assignment, the “initialization” routine for the interface is called.
<b>SIOCGIFADDR</b>	Get interface address.
<b>SIOCSIFDSTADDR</b>	Set point to point address for interface.
<b>SIOCGIFDSTADDR</b>	Get point to point address for interface.
<b>SIOCSIFFLAGS</b>	Set interface flags field. If the interface is marked down, any processes currently routing packets through the interface are notified.
<b>SIOCGIFFLAGS</b>	Get interface flags.
<b>SIOCGIFCONF</b>	Get interface configuration list. This request takes an <code>ifconf</code> structure (see below) as a value-result parameter. The <code>ifc_len</code> field should be initially set to the size of the buffer pointed to by <code>ifc_buf</code> . On return it will contain the length, in bytes, of the configuration list.

```

/*
 * Structure used in SIOCGIFCONF request.
 * Used to retrieve interface configuration
 * for machine (useful for programs which
 * must know all networks accessible).
 */
struct ifconf {
    int    ifc_len;        /* size of associated buffer */
    union {
        caddr_t ifcu_buf;
        struct ifreq *ifcu_req;
    } ifc_ifcu;
#define ifc_buf ifc_ifcu.ifcu_buf /* buffer address */
#define ifc_req ifc_ifcu.ifcu_req /* array of structures returned */
};

```

**SIOCADDMULTI** Enable a multicast address for the interface. A maximum of 64 multicast addresses may be enabled for any given interface.

**SIOCDELMULTI** Disable a previously set multicast address.

**SIOCSMISC** Toggle promiscuous mode.

**SEE ALSO**

**arp(4P), lo(4)**

**NAME**

inet – Internet protocol family

**SYNOPSIS**

options INET

#include <sys/types.h>

#include <netinet/in.h>

**DESCRIPTION**

The Internet protocol family implements a collection of protocols which are centered around the *Internet Protocol* (IP) and which share a common address format. The Internet family provides protocol support for the `SOCK_STREAM`, `SOCK_DGRAM`, and `SOCK_RAW` socket types.

**PROTOCOLS**

The Internet protocol family is comprised of the Internet Protocol (IP), the Address Resolution Protocol (ARP), the Internet Control Message Protocol (ICMP), the Transmission Control Protocol (TCP), and the User Datagram Protocol (UDP).

TCP is used to support the `SOCK_STREAM` abstraction while UDP is used to support the `SOCK_DGRAM` abstraction; see `tcp(4P)` and `udp(4P)`. A raw interface to IP is available by creating an Internet socket of type `SOCK_RAW`; see `ip(4P)`. ICMP is used by the kernel to handle and report errors in protocol processing. It is also accessible to user programs; see `icmp(4P)`. ARP is used to translate 32-bit IP addresses into 48-bit Ethernet addresses; see `arp(4P)`.

The 32-bit IP address is divided into network number and host number parts. It is frequency-encoded; the most-significant bit is zero in Class A addresses, in which the high-order 8 bits are the network number. Class B addresses have their high order two bits set to 10 and use the high-order 16 bits as the network number field. Class C addresses have a 24-bit network number part of which the high order three bits are 110. Sites with a cluster of local networks may chose to use a single network number for the cluster; this is done by using subnet addressing. The local (host) portion of the address is further subdivided into subnet number and host number parts. Within a subnet, each subnet appears to be an individual network; externally, the entire cluster appears to be a single, uniform network requiring only a single routing entry. Subnet addressing is enabled and examined by the following `ioctl(2)` commands on a datagram socket in the Internet domain; they have the same form as the `SIOCIFADDR` command (see `intro(4)`).

**SIOCSIFNETMASK** Set interface network mask. The network mask defines the network part of the address; if it contains more of the address than the address type would indicate, then subnets are in use.

**SIOCGIFNETMASK** Get interface network mask.

**ADDRESSING**

IP addresses are four byte quantities, stored in network byte order (on Sun386i systems these are word and byte reversed).

Sockets in the Internet protocol family use the following addressing structure:

```
struct sockaddr_in {
    short    sin_family;
    u_short sin_port;
    struct   in_addr sin_addr;
    char     sin_zero[8];
};
```

Library routines are provided to manipulate structures of this form; see `intro(3)`.

The `sin_addr` field of the `sockaddr_in` structure specifies a local or remote IP address. Each network interface has its own unique IP address. The special value `INADDR_ANY` may be used in this field to effect “wildcard” matching. Given in a `bind(2)` call, this value leaves the local IP address of the socket unspecified, so that the socket will receive connections or messages directed at any of the valid IP addresses of the system. This can prove useful when a process neither knows nor cares what the local IP

address is or when a process wishes to receive requests using all of its network interfaces. The `sockaddr_in` structure given in the `bind(2)` call must specify an `in_addr` value of either `IPADDR_ANY` or one of the system's valid IP addresses. Requests to bind any other address will elicit the error `EADDRNOTAVAIL`. When a `connect(2)` call is made for a socket that has a wildcard local address, the system sets the `sin_addr` field of the socket to the IP address of the network interface that the packets for that connection are routed via.

The `sin_port` field of the `sockaddr_in` structure specifies a port number used by TCP or UDP. The local port address specified in a `bind(2)` call is restricted to be greater than `IPPORT_RESERVED` (defined in `<netinet/in.h>`) unless the creating process is running as the super-user, providing a space of protected port numbers. In addition, the local port address must not be in use by any socket of same address family and type. Requests to bind sockets to port numbers being used by other sockets return the error `EADDRINUSE`. If the local port address is specified as 0, then the system picks a unique port address greater than `IPPORT_RESERVED`. A unique local port address is also picked when a socket which is not bound is used in a `connect(2)` or `send(2)` call. This allows programs which do not care which local port number is used to set up TCP connections by simply calling `socket(2)` and then `connect(2)`, and to send UDP datagrams with a `socket(2)` call followed by a `send(2)` call.

Although this implementation restricts sockets to unique local port numbers, TCP allows multiple simultaneous connections involving the same local port number so long as the remote IP addresses or port numbers are different for each connection. Programs may explicitly override the socket restriction by setting the `SO_REUSEADDR` socket option with `setsockopt` (see `getsockopt(2)`).

#### SEE ALSO

`bind(2)`, `connect(2)`, `getsockopt(2)`, `ioctl(2)`, `send(2)`, `socket(2)`, `intro(3)`, `byteorder(3N)`, `gethostent(3N)`, `getnetent(3N)`, `getprotoent(3N)`, `getservent(3N)`, `inet(3N)`, `intro(4)`, `arp(4P)`, `icmp(4P)`, `ip(4P)` `tcp(4P)`, `udp(4P)`

Network Information Center, *DDN Protocol Handbook* (3 vols.), Network Information Center, SRI International, Menlo Park, Calif., 1985.

*A 4.2BSD Interprocess Communication Primer*

#### WARNING

The Internet protocol support is subject to change as the Internet protocols develop. Users should not depend on details of the current implementation, but rather the services exported.



## NAME

ip – Internet Protocol

## SYNOPSIS

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

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

```
s = socket(AF_INET, SOCK_RAW, proto);
```

## DESCRIPTION

IP is the internetwork datagram delivery protocol that is central to the Internet protocol family. Programs may use IP through higher-level protocols such as the Transmission Control Protocol (TCP) or the User Datagram Protocol (UDP), or may interface directly using a “raw socket.” See [tcp\(4P\)](#) and [udp\(4P\)](#). The protocol options defined in the IP specification may be set in outgoing datagrams.

Raw IP sockets are connectionless and are normally used with the `sendto` and `recvfrom` calls, (see `send(2)` and `recv(2)`) although the `connect(2)` call may also be used to fix the destination for future datagrams (in which case the `read(2V)` or `recv(2)` and `write(2V)` or `send(2)` calls may be used). If `proto` is zero, the default protocol, `IPPROTO_RAW`, is used. If `proto` is non-zero, that protocol number will be set in outgoing datagrams and will be used to filter incoming datagrams. An IP header will be generated and prepended to each outgoing datagram; Received datagrams are returned with the IP header and options intact.

A single socket option, `IP_OPTIONS`, is supported at the IP level. This socket option may be used to set IP options to be included in each outgoing datagram. IP options to be sent are set with `setsockopt` (see `getsockopt(2)`). The `getsockopt(2)` call returns the IP options set in the last `setsockopt` call. IP options on received datagrams are visible to user programs only using raw IP sockets. The format of IP options given in `setsockopt` matches those defined in the IP specification with one exception: the list of addresses for the source routing options must include the first-hop gateway at the beginning of the list of gateways. The first-hop gateway address will be extracted from the option list and the size adjusted accordingly before use. IP options may be used with any socket type in the Internet family.

At the socket level, the socket option `SO_DONTROUTE` may be applied. This option forces datagrams being sent to bypass the routing step in output. Normally, IP selects a network interface to send the datagram via, and possibly an intermediate gateway, based on an entry in the routing table. See [routing\(4N\)](#). When `SO_DONTROUTE` is set, the datagram will be sent via the interface whose network number or full IP address matches the destination address. If no interface matches, the error `ENETUNRCH` will be returned.

Datagrams flow through the IP layer in two directions: from the network `ip` to user processes and from user processes *down* to the network. Using this orientation, IP is layered *above* the network interface drivers and *below* the transport protocols such as UDP and TCP. The Internet Control Message Protocol (ICMP) is logically a part of IP. See [icmp\(4P\)](#).

IP provides for a checksum of the header part, but not the data part of the datagram. The checksum value is computed and set in the process of sending datagrams and checked when receiving datagrams. IP header checksumming may be disabled for debugging purposes by patching the kernel variable `ipcksum` to have the value zero.

IP options in received datagrams are processed in the IP layer according to the protocol specification. Currently recognized IP options include: security, loose source and record route (LSRR), strict source and record route (SSRR), record route, stream identifier, and internet timestamp.

The IP layer will normally forward received datagrams that are not addressed to it. Forwarding is under the control of the kernel variable `ipforwarding`: if `ipforwarding` is zero, IP datagrams will not be forwarded; if `ipforwarding` is one, IP datagrams will be forwarded. `ipforwarding` is usually set to one only in machines with more than one network interface (internetwork routers). This kernel variable can be patched to enable or disable forwarding.

The IP layer will send an ICMP message back to the source host in many cases when it receives a datagram that can not be handled. A "time exceeded" ICMP message will be sent if the "time to live" field in the IP header drops to zero in the process of forwarding a datagram. A "destination unreachable" message will be sent if a datagram can not be forwarded because there is no route to the final destination, or if it can not be fragmented. If the datagram is addressed to the local host but is destined for a protocol that is not supported or a port that is not in use, a destination unreachable message will also be sent. The IP layer may send an ICMP "source quench" message if it is receiving datagrams too quickly. ICMP messages are only sent for the first fragment of a fragmented datagram and are never returned in response to errors in other ICMP messages.

The IP layer supports fragmentation and reassembly. Datagrams are fragmented on output if the datagram is larger than the maximum transmission unit (MTU) of the network interface. Fragments of received datagrams are dropped from the reassembly queues if the complete datagram is not reconstructed within a short time period.

Errors in sending discovered at the network interface driver layer are passed by IP back up to the user process.

## ERRORS

A socket operation may fail with one of the following errors returned:

EACCESS	when specifying an IP broadcast destination address if the caller is not the super-user;
EISCONN	when trying to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected;
EMSGSIZE	when sending datagram that is too large for an interface, but is not allowed be fragmented (such as broadcasts);
ENETUNREACH	when trying to establish a connection or send a datagram, if there is no matching entry in the routing table, or if an ICMP "destination unreachable" message is received.
ENOTCONN	when trying to send a datagram, but no destination address is specified, and the socket hasn't been connected;
ENOBUFS	when the system runs out of memory for fragmentation buffers or other internal data structure;
EADDRNOTAVAIL	when an attempt is made to create a socket with a local address that matches no network interface, or when specifying an IP broadcast destination address and the network interface does not support broadcast;

The following errors may occur when setting or getting IP options:

EINVAL	An unknown socket option name was given.
EINVAL	The IP option field was improperly formed; an option field was shorter than the minimum value or longer than the option buffer provided.

## SEE ALSO

**connect(2), getsockopt(2), read(2V), recv(2), send(2), write(2V), icmp(4P), inet(4F) routing(4N), tcp(4P), udp(4P)**

Postel, Jon, "*Internet Protocol - DARPA Internet Program Protocol Specification*," RFC 791, Network Information Center, SRI International, Menlo Park, Calif., September 1981. (Sun 800-1063-01)

**BUGS**

Raw sockets should receive ICMP error packets relating to the protocol; currently such packets are simply discarded.

Users of higher-level protocols such as TCP and UDP should be able to see received IP options.

## NAME

kb – Sun keyboard STREAMS module

## CONFIG

pseudo-device *kbnumber*

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/stream.h>
#include <sys/stropts.h>
#include <sundev/vuid_event.h>
#include <sundev/kbio.h>
#include <sundev/kbd.h>

ioctl(fd, I_PUSH, "kb");
```

## DESCRIPTION

The **kb** STREAMS module processes byte streams generated by Sun keyboards attached to a CPU serial or parallel port. Definitions for altering keyboard translation, and reading events from the keyboard, are in `<sundev/kbio.h>` and `<sundev/kbd.h>`. *number* specifies the maximum number of keyboards supported by the system.

**kb** recognizes which keys have been typed using a set of tables for each known type of keyboard. Each translation table is an array of 128 16-bit words (**unsigned shorts**). If an entry in the table is less than 0x100, it is treated as an ISO 8859/1 character. Higher values indicate special characters that invoke more complicated actions.

## Keyboard Translation Mode

The keyboard can be in one of the following translation modes:

TR_NONE	Keyboard translation is turned off and up/down key codes are reported.
TR_ASCII	ISO 8859/1 codes are reported.
TR_EVENT	<b>firm_events</b> are reported (see <i>SunView Programmer's Guide</i> ).
TR_UNTRANS_EVENT	<b>firm_events</b> containing unencoded keystation codes are reported for all input events within the window system.

## Keyboard Translation-Table Entries

All instances of the **kb** module share seven translation tables used to convert raw keystation codes to event values. The tables are:

Unshifted	Used when a key is depressed and no shifts are in effect.
Shifted	Used when a key is depressed and a Shift key is being held down.
Caps Lock	Used when a key is depressed and Caps Lock is in effect.
Alt Graph	Used when a key is depressed and the Alt Graph key is being held down.
Num Lock	Used when a key is depressed and Num Lock is in effect.
Controlled	Used when a key is depressed and the Control key is being held down (regardless of whether a Shift key or the Alt Graph is being held down, or whether Caps Lock or Num Lock is in effect).
Key Up	Used when a key is released.

Each key on the keyboard has a "key station" code which is a number from 0 to 127. This number is used as an index into the translation table that is currently in effect. If the corresponding entry in that translation table is a value from 0 to 255, this value is treated as an ISO 8859/1 character, and that character is the result of the translation.

If the entry is a value above 255, it is a “special” entry. Special entry values are classified according to the value of the high-order bits. The high-order value for each class is defined as a constant, as shown in the list below. The value of the low-order bits, when added to this constant, distinguishes between keys within each class:

SHIFTKEYS 0x100	A shift key. The value of the particular shift key is added to determine which shift mask to apply:
CAPSLock 0	“Caps Lock” key.
SHIFTLOCK 1	“Shift Lock” key.
LEFTSHIFT 2	Left-hand “Shift” key.
RIGHTSHIFT 3	Right-hand “Shift” key.
LEFTCTRL 4	Left-hand (or only) “Control” key.
RIGHTCTRL 5	Right-hand “Control” key.
ALTGRAPH 9	“Alt Graph” key.
ALT 10	“Alternate” key on the Sun-3 keyboard, or “Alt” key on the Sun-4 keyboard.
NUMLOCK 11	“Num Lock” key.
BUCKYBITS 0x200	Used to toggle mode-key-up/down status without altering the value of an accompanying ISO 8859/1 character. The actual bit-position value, minus 7, is added.
METABIT 0	The “Meta” key was pressed along with the key. This is the only user-accessible bucky bit. It is ORed in as the 0x80 bit; since this bit is a legitimate bit in a character, the only way to distinguish between, for example, 0xA0 as META+0x20 and 0xA0 as an 8-bit character is to watch for “META key up” and “META key down” events and keep track of whether the META key was down.
SYSTEMBIT 1	The “System” key was pressed. This is a place holder to indicate which key is the system-abort key.
FUNNY 0x300	Performs various functions depending on the value of the low 4 bits:
NOP 0x300	Does nothing.
OOPS 0x301	Exists, but is undefined.
HOLE 0x302	There is no key in this position on the keyboard, and the position-code should not be used.
NOSCROLL 0x303	Alternately sends CTRL-S and CTRL-Q characters.
CTRLS 0x304	Sends CTRL-S character and toggles NOScroll key.
CTRLQ 0x305	Sends CTRL-Q character and toggles NOScroll key.
RESET 0x306	Keyboard reset.
ERROR 0x307	The keyboard driver detected an internal error.
IDLE 0x308	The keyboard is idle (no keys down).
COMPOSE 0x309	This key is the COMPOSE key; the next two keys should comprise a two-character “COMPOSE key” sequence.

	NONL 0x30A	Used only in the Num Lock table; indicates that this key is not affected by the Num Lock state, so that the translation table to use to translate this key should be the one that would have been used had Num Lock not been in effect.
	0x30B — 0x30F	Reserved for nonparameterized functions.
FA_CLASS 0x400		This key is a “floating accent” or “dead” key. When this key is pressed, the next key generates an event for an accented character; for example, “floating accent grave” followed by the “a” key generates an event with the ISO 8859/1 code for the “a with grave accent” character. The low-order bits indicate which accent; the codes for the individual “floating accents” are as follows:
	FA_UMLAUT 0x400	umlaut
	FA_CFLEX 0x401	circumflex
	FA_TILDE 0x402	tilde
	FA_CEDILLA 0x403	cedilla
	FA_ACUTE 0x404	acute accent
	FA_GRAVE 0x405	grave accent
STRING 0x500		The low-order bits index a table of strings. When a key with a STRING entry is depressed, the characters in the null-terminated string for that key are sent, character by character. The maximum length is defined as:
	KTAB_STRLEN 10	
		Individual string numbers are defined as:
	HOMEARROW 0x00	
	UPARROW 0x01	
	DOWNARROW 0x02	
	LEFTARROW 0x03	
	RIGHTARROW 0x04	
		String numbers 0x05 — 0x0F are available for custom entries.
FUNCKEYS 0x600		Function keys. The next-to-lowest 4 bits indicate the group of function keys:
	LEFTFUNC 0x600	
	RIGHTFUNC 0x610	
	TOPFUNC 0x620	
	BOTTOMFUNC 0x630	
		The low 4 bits indicate the function key number within the group:
	LF( <i>n</i> )	(LEFTFUNC+( <i>n</i> )-1)
	RF( <i>n</i> )	(RIGHTFUNC+( <i>n</i> )-1)
	TF( <i>n</i> )	(TOPFUNC+( <i>n</i> )-1)
	BF( <i>n</i> )	(BOTTOMFUNC+( <i>n</i> )-1)
		There are 64 keys reserved for function keys. The actual positions may not be on left/right/top/bottom of the keyboard, although they usually are.
PADKEYS 0x700		This key is a “numeric keypad key.” These entries should appear only in the Num Lock translation table; when Num Lock is in effect, these events will be generated by pressing keys on the right-hand keypad. The low-order bits indicate which key; the codes for the individual keys are as follows:

PADEQUAL 0x700	“=” key
PADSLASH 0x701	“/” key
PADSTAR 0x702	“*” key
PADMINUS 0x703	“-” key
PADSEP 0x704	“,” key
PAD7 0x705	“7” key
PAD8 0x706	“8” key
PAD9 0x707	“9” key
PADPLUS 0x708	“+” key
PAD4 0x709	“4” key
PAD5 0x70A	“5” key
PAD6 0x70B	“6” key
PAD1 0x70C	“1” key
PAD2 0x70D	“2” key
PAD3 0x70E	“3” key
PAD0 0x70F	“0” key
PADDOT 0x710	“.” key
PADENTER 0x711	“Enter” key

In `TR_ASCII` mode, when a function key is pressed, the following escape sequence is sent:

```
ESC[0...9z
```

where `ESC` is a single escape character and “0..9” indicates the decimal representation of the function-key value. For example, function key `R1` sends the sequence:

```
ESC[208z
```

because the decimal value of `RF(1)` is 208. In `TR_EVENT` mode, if there is a `VUID` event code for the function key in question, an event with that event code is generated; otherwise, individual events for the characters of the escape sequence are generated.

#### Keyboard Compatibility Mode

`kb` is in “compatibility mode” when it starts up. In this mode, when the keyboard is in the `TR_EVENT` translation mode, ISO 8859/1 characters from the “upper half” of the character set (that is, characters with the 8th bit set) are presented as events with codes in the `ISO_FIRST` range (as defined in `<sundev/vuid_event.h>`). The event code is `ISO_FIRST` plus the character value. This is for backwards compatibility with older versions of the keyboard driver. If compatibility mode is turned off, ISO 8859/1 characters are presented as events with codes equal to the character code.

#### IOCTLS

The following `ioctl()` requests set and retrieve the current translation mode of a keyboard:

**KIOCTRANS** The argument is a pointer to an `int`. The translation mode is set to the value in the `int` pointed to by the argument.

**KIOCGTRANS** The argument is a pointer to an `int`. The current translation mode is stored in the `int` pointed to by the argument.

`ioctl()` requests for changing and retrieving entries from the keyboard translation table use the `kiockeymap` structure:

```

struct kiockeymap {
    int    kio_tablemask; /* Translation table (one of: 0, CAPSMASK,
                          SHIFTMASK, CTRLMASK, UPMASK,
                          ALTGRAPHMASK, NUMLOCKMASK) */
#define KIOABORT1  -1 /* Special "mask": abort1 keystation */
#define KIOABORT2  -2 /* Special "mask": abort2 keystation */
    u_char kio_station; /* Physical keyboard key station (0-127) */
    u_short kio_entry; /* Translation table station's entry */
    char   kio_string[10]; /* Value for STRING entries (null terminated) */
};

```

**KIOCSKEY** The argument is a pointer to a **kiockeymap** structure. The translation table entry referred to by the values in that structure is changed.

**kio\_tablemask** specifies which of the five translation tables contains the entry to be modified:

UPMASK 0x0080	“Key Up” translation table.
NUMLOCKMASK 0x0800	“Num Lock” translation table.
CTRLMASK 0x0030	“Controlled” translation table.
ALTGRAPHMASK 0x0200	“Alt Graph” translation table.
SHIFTMASK 0x000E	“Shifted” translation table.
CAPSMASK 0x0001	“Caps Lock” translation table.
(No shift keys pressed or locked)	“Unshifted” translation table.

**kio\_station** specifies the keystation code for the entry to be modified. The value of **kio\_entry** is stored in the entry in question. If **kio\_entry** is between **STRING** and **STRING+15**, the string contained in **kio\_string** is copied to the appropriate string table entry. This call may return **EINVAL** if there are invalid arguments.

There are a couple special values of **kio\_tablemask** that affect the two step “break to the PROM monitor” sequence. The usual sequence is **SETUP-a** or **L1-a**. If **kio\_tablemask** is **KIOABORT1** then the value of **kio\_station** is set to be the first keystation in the sequence. If **kio\_tablemask** is **KIOABORT2** then the value of **kio\_station** is set to be the second keystation in the sequence.

**KIOCGKEY** The argument is a pointer to a **kiockeymap** structure. The current value of the keyboard translation table entry specified by **kio\_tablemask** and **kio\_station** is stored in the structure pointed to by the argument. This call may return **EINVAL** if there are invalid arguments.

**KIOCTYPE** The argument is a pointer to an **int**. A code indicating the type of the keyboard is stored in the **int** pointed to by the argument:

<b>KB_KLUNK</b>	Micro Switch 103SD32-2
<b>KB_VT100</b>	Keytronics VT100 compatible
<b>KB_SUN2</b>	Sun-2 keyboard
<b>KB_SUN3</b>	Sun-3 keyboard
<b>KB_SUN4</b>	Sun-4 keyboard
<b>KB_ASCII</b>	ASCII terminal masquerading as keyboard

-1 is stored in the **int** pointed to by the argument if the keyboard type is unknown.

**KIOCLAYOUT** The argument is a pointer to an **int**. On a Sun-4 keyboard, the layout code specified by the keyboard’s DIP switches is stored in the **int** pointed to by the argument.



- KIOCCMD** The argument is a pointer to an `int`. The command specified by the value of the `int` pointed to by the argument is sent to the keyboard. The commands that can be sent are:
- Commands to the Sun-2, Sun-3, and Sun-4 keyboard:
- KBD\_CMD\_RESET** Reset keyboard as if power-up.
  - KBD\_CMD\_BELL** Turn on the bell.
  - KBD\_CMD\_NOBELL** Turn off the bell
- Commands to the Sun-3 and Sun-4 keyboard:
- KBD\_CMD\_CLICK** Turn on the click annunciator.
  - KBD\_CMD\_NOCLICK** Turn off the click annunciator.
- Inappropriate commands for particular keyboard types are ignored. Since there is no reliable way to get the state of the bell or click (because we cannot query the keyboard, and also because a process could do writes to the appropriate serial driver — thus going around this `ioctl()` request) we do not provide an equivalent `ioctl()` to query its state.
- KIOCSLED** The argument is a pointer to an `char`. On the Sun-4 keyboard, the LEDs are set to the value specified in that `char`. The values for the four LEDs are:
- LED\_CAPS\_LOCK** “Caps Lock” light.
  - LED\_COMPOSE** “Compose” light.
  - LED\_SCROLL\_LOCK** “Scroll Lock” light.
  - LED\_NUM\_LOCK** “Num Lock” light.
- KIOCGLED** The argument is a pointer to a `char`. The current state of the LEDs is stored in the `char` pointed to by the argument.
- KIOCSCOMPAT** The argument is a pointer to an `int`. “Compatibility mode” is turned on if the `int` has a value of 1, and is turned off if the `int` has a value of 0.
- KIOCGCOMPAT** The argument is a pointer to an `int`. The current state of “compatibility mode” is stored in the `int` pointed to by the argument.
- KIOCGDIRECT** These `ioctl()` requests are supported for compatibility with the system keyboard device `/dev/kbd`. **KIOCSDIRECT** has no effect, and **KIOCGDIRECT** always returns 1.

**SEE ALSO**

`click(1)`, `loadkeys(1)`, `kbd(4S)`, `termio(4)`, `win(4S)`, `keytables(5)`

*SunView Programmer's Guide* (describes `firm_event` format)

**NAME**

kbd – Sun keyboard

**CONFIG**

None; included in standard system.

**DESCRIPTION**

The **kbd** device provides access to the Sun Workstation keyboard. When opened, it provides access to the standard keyboard device for the workstation (attached either to a CPU serial or parallel port). It is a multiplexing driver; a stream referring to the standard keyboard device, with the **kb(4M)** and **ttcompat(4M)** STREAMS modules pushed on top of that device, is linked below it. Normally, this device passes input to the “workstation console” driver, which is linked above a special minor device of **kbd**, so that keystrokes appear as input on **/dev/console**; the **KIOCSDIRECT ioctl** must be used to direct input towards or away from the **/dev/kbd** device.

**IOCTLS**

**KIOCSDIRECT** The argument is a pointer to an **int**. If the value in the **int** pointed to by the argument is 1, subsequent keystrokes typed on the system keyboard will be sent to **/dev/kbd**; if it is 0, subsequent keystrokes will be sent to the “workstation console” device. When the last process that has **/dev/kbd** open closes it, if keystrokes had been sent to **/dev/kbd** they are redirected back to the “workstation console” device.

**KIOCGDIRECT** The argument is a pointer to an **int**. If keystrokes are currently being sent to **/dev/kbd**, 1 is stored in the **int** pointed to by the argument; if keystrokes are currently being sent to the “workstation console” device, 0 is stored there.

**FILES**

**/dev/kbd**

**SEE ALSO**

**console(4S)**, **kb(4M)**, **ttcompat(4M)**, **win(4S)**, **zs(4S)**

**NAME**

ldterm – standard terminal STREAMS module

**CONFIG**

None; included by default.

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/stream.h>
#include <sys/stropts.h>
```

```
ioctl(fd, I_PUSH, "ldterm");
```

**DESCRIPTION**

**ldterm** is a STREAMS module that provides most of the **termio(4)** terminal interface. This module does not perform the low-level device control functions specified by flags in the **c\_cflag** word of the **termios** structure or by the **IGNBRK**, **IGNPAR**, **PARMRK**, or **INPCK** flags in the **c\_iflag** word of the **termios** structure; those functions must be performed by the driver or by modules pushed below the **ldterm** module. All other **termio** functions are performed by **ldterm**; some of them, however, require the cooperation of the driver or modules pushed below **ldterm**, and may not be performed in some cases. These include the **IXOFF** flag in the **c\_iflag** word and the delays specified in the **c\_oflag** word.

**Read-side Behavior**

Various types of STREAMS messages are processed as follows:

- M\_BREAK** When this message is received, either an interrupt signal is generated, or the message is treated as if it were an **M\_DATA** message containing a single ASCII NUL character, depending on the state of the **BRKINT** flag.
- M\_DATA** These messages are normally processed using the standard **termio** input processing. If the **ICANON** flag is set, a single input record (“line”) is accumulated in an internal buffer, and sent upstream when a line-terminating character is received. If the **ICANON** flag is not set, other input processing is performed and the processed data is passed upstream.
- If output is to be stopped or started as a result of the arrival of characters, **M\_STOP** and **M\_START** messages are sent downstream, respectively. If the **IXOFF** flag is set, and input is to be stopped or started as a result of flow-control considerations, **M\_STOPI** and **M\_STARTI** messages are sent downstream, respectively.
- M\_DATA** messages are sent downstream, as necessary, to perform echoing.
- If a signal is to be generated, a **M\_FLUSH** message with a flag byte of **FLUSHR** is placed on the read queue, and if the signal is also to flush output a **M\_FLUSH** message with a flag byte of **FLUSHW** is sent downstream.
- M\_CTL** If the first byte of the message is **MC\_NOCANON**, the input processing normally performed on **M\_DATA** messages is disabled, and those messages are passed upstream unmodified; this is for the use of modules or drivers that perform their own input processing, such as a pseudo-terminal in **TIOCREMOTE** mode connected to a program that performs this processing. If the first byte of the message is **MC\_DOCANON**, the input processing is enabled. Otherwise, the message is ignored; in any case, the message is passed upstream.
- M\_FLUSH** The read queue of the module is flushed of all its data messages, and all data in the record being accumulated is also flushed. The message is passed upstream.
- M\_HANGUP** Data is flushed as it is for a **M\_FLUSH** message, and **M\_FLUSH** messages with a flag byte of **FLUSHRW** are sent upstream and downstream. Then an **M\_PCSIG** message is sent upstream with a signal of **SIGCONT**, followed by the **M\_HANGUP** message.
- M\_IOCACK** The data contained within the message, which is to be returned to the process, is augmented if necessary, and the message is passed upstream.

All other messages are passed upstream unchanged.

#### Write-side behavior

Various types of STREAMS messages are processed as follows:

- M\_FLUSH** The write queue of the module is flushed of all its data messages, and the message is passed downstream.
- M\_IOCTL** The function to be performed for this `ioctl()` request by the `ldterm` module is performed, and the message is passed downstream in most cases. The `TCFLSH` and `TCXONC` `ioctl()` requests can be performed entirely in this module, so the reply is sent upstream and the message is not passed downstream.
- M\_DATA** If the `OPOST` flag is set, or both the `XCASE` and `ICANON` flags are set, output processing is performed and the processed message is passed downstream, along with any `M_DELAY` messages generated. Otherwise, the message is passed downstream without change.

All other messages are passed downstream unchanged.

#### IOCTLS

The following `ioctl()` requests are processed by the `ldterm` module. All others are passed downstream.

#### TCGETS

- TCGETA** The message is passed downstream; if an acknowledgment is seen, the data provided by the driver and modules downstream is augmented and the acknowledgement is passed upstream.

#### TCSETS

#### TCSETSW

#### TCSETSF

#### TCSETA

#### TCSETAW

#### TCSETAF

The parameters that control the behavior of the `ldterm` module are changed. If a mode change requires options at the stream head to be changed, a `M_SETOPT` message is sent upstream. If the `ICANON` flag is turned on or off, the read mode at the stream head is changed to message-nondiscard or byte-stream mode, respectively. If it is turned on, the `vmin` and `vtime` values at the stream head are set to 1 and 0, respectively; if it is turned off, they are set to the values specified by the `ioctl()` request. The `vmin` and `vtime` values are also set if `ICANON` is off and the values are changed by the `ioctl()` request. If the `TOSTOP` flag is turned on or off, the `tostop` mode at the stream head is turned on or off, respectively.

#### TCFLSH

If the argument is 0, an `M_FLUSH` message with a flag byte of `FLUSHR` is sent downstream and placed on the read queue. If the argument is 1, the write queue is flushed of all its data messages and a `M_FLUSH` message with a flag byte of `FLUSHW` is sent upstream and downstream. If the argument is 2, the write queue is flushed of all its data messages and a `M_FLUSH` message with a flag byte of `FLUSHRW` is sent downstream and placed on the read queue.

#### TCXONC

If the argument is 0, and output is not already stopped, an `M_STOP` message is sent downstream. If the argument is 1, and output is stopped, an `M_START` message is sent downstream. If the argument is 2, and input is not already stopped, an `M_STOPI` message is sent downstream. If the argument is 3, and input is stopped, an `M_STARTI` message is sent downstream.

#### SEE ALSO

`console(4S)`, `mcp(4S)`, `mti(4S)`, `pty(4)`, `termio(4)`, `ttcompat(4M)`, `zs(4S)`

**NAME**

ie – Intel 10 Mb/s Ethernet interface

**CONFIG — SUN-4 SYSTEM**

device ie0 at obio ? csr 0xf6000000 priority 3  
 device ie1 at vme24d16 ? csr 0xe88000 priority 3 vector ieintr 0x75  
 device ie2 at vme24d16 ? csr 0x31ff02 priority 3 vector ieintr 0x76  
 device ie3 at vme24d16 ? csr 0x35ff02 priority 3 vector ieintr 0x77

**CONFIG — SUN-3x SYSTEM**

device ie0 at obio ? csr 0x65000000 priority 3  
 device ie1 at vme24d16 ? csr 0xe88000 priority 3 vector ieintr 0x75

**CONFIG — SUN-3 SYSTEM**

device ie0 at obio ? csr 0xc0000 priority 3  
 device ie1 at vme24d16 ? csr 0xe88000 priority 3 vector ieintr 0x75  
 device ie2 at vme24d32 ? csr 0x31ff02 priority 3 vector ieintr 0x76  
 device ie3 at vme24d32 ? csr 0x35ff02 priority 3 vector ieintr 0x77

**CONFIG — SUN-3E SYSTEM**

device ie0 at vme24d16 ? csr 0x31ff02 priority 3 vector ieintr 0x74

**CONFIG — SUN386i SYSTEM**

device ie0 at obmem ? csr 0xD0000000 irq 21 priority 3

**DESCRIPTION**

The ie interface provides access to a 10 Mb/s Ethernet network through a controller using the Intel 82586 LAN Coprocessor chip. For a general description of network interfaces see if(4N).

ie0 specifies a CPU-board-resident interface, except on a Sun-3E where ie0 is the Sun-3/E Ethernet expansion board. ie1 specifies a Multibus Intel Ethernet interface for use with a VME adapter. ie2 and ie3 specify SunNet Ethernet/VME Controllers, also known as a Sun-3/E Ethernet expansion boards.

**SEE ALSO**

if(4N), ie(4S)

**DIAGNOSTICS**

There are too many driver messages to list them all individually here. Some of the more common messages and their meanings follow.

**ie%d: Ethernet jammed**

Network activity has become so intense that sixteen successive transmission attempts failed, and the 82586 gave up on the current packet. Another possible cause of this message is a noise source somewhere in the network, such as a loose transceiver connection.

**ie%d: no carrier**

The 82586 has lost input to its carrier detect pin while trying to transmit a packet, causing the packet to be dropped. Possible causes include an open circuit somewhere in the network and noise on the carrier detect line from the transceiver.

**ie%d: lost interrupt: resetting**

The driver and 82586 chip have lost synchronization with each other. The driver recovers by resetting itself and the chip.

**ie%d: iebark reset**

The 82586 failed to complete a watchdog timeout command in the allotted time. The driver recovers by resetting itself and the chip.

**ie%d: WARNING: requeuing**

The driver has run out of resources while getting a packet ready to transmit. The packet is put back on the output queue for retransmission after more resources become available.

**ie%d: panic: scb overwritten**

The driver has discovered that memory that should remain unchanged after initialization has become corrupted. This error usually is a symptom of a bad 82586 chip.

**ie%d: giant packet**

Provided that all stations on the Ethernet are operating according to the Ethernet specification, this error "should never happen," since the driver allocates its receive buffers to be large enough to hold packets of the largest permitted size. The most likely cause of this message is that some other station on the net is transmitting packets whose lengths exceed the maximum permitted for Ethernet.

**NAME**

lo – software loopback network interface

**SYNOPSIS**

**pseudo-device loop**

**DESCRIPTION**

The **loop** device is a software loopback network interface; see **if(4N)** for a general description of network interfaces.

The **loop** interface is used for performance analysis and software testing, and to provide guaranteed access to Internet protocols on machines with no local network interfaces. A typical application is the **comsat(8C)** server which accepts notification of mail delivery through a particular port on the loopback interface.

By default, the loopback interface is accessible at Internet address 127.0.0.1 (non-standard); this address may be changed with the **SIOCSIFADDR** ioctl.

**SEE ALSO**

**if(4N)**, **inet(4F)**, **comsat(8C)**

**DIAGNOSTICS**

**lo%d: can't handle af%d**

The interface was handed a message with addresses formatted in an unsuitable address family; the packet was dropped.

**BUGS**

It should handle all address and protocol families. An approved network address should be reserved for this interface.

**NAME**

lofs – loopback virtual file system

**CONFIG**

options LOFS

**SYNOPSIS**

```
#include <sys/mount.h>
mount(MOUNT_LOFS, virtual, flags, dir);
```

**DESCRIPTION**

The loopback file system device allows new, virtual file systems to be created, which provide access to existing files using alternate pathnames. Once the virtual file system is created, other file systems can be mounted within it without affecting the original file system. File systems that are subsequently mounted onto the original file system, however, *are* visible to the virtual file system, unless or until the corresponding mount point in the virtual file system is covered by a file system mounted there.

*virtual* is the mount point for the virtual file system. *dir* is the pathname of the existing file system. *flags* is either 0 or `M_RDONLY`. The `M_RDONLY` flag forces all accesses in the new name space to be read-only; without it, accesses are the same as for the underlying file system. All other `mount(2V)` flags are preserved from the underlying file systems.

A loopback mount of '/' onto `/tmp/newroot` allows the entire file system hierarchy to appear as if it were duplicated under `/tmp/newroot`, including any file systems mounted from remote NFS servers. All files would then be accessible either from a pathname relative to '/', or from a pathname relative to `/tmp/newroot` until such time as a file system is mounted in `/tmp/newroot`, or any of its subdirectories.

Loopback mounts of '/' can be performed in conjunction with the `chroot(2)` system call, to provide a complete virtual file system to a process or family of processes.

Recursive traversal of loopback mount points is not allowed; after the loopback mount of `/tmp/newroot`, the file `/tmp/newroot/tmp/newroot` does not contain yet another file system hierarchy; rather, it appears just as `/tmp/newroot` did before the loopback mount was performed (say, as an empty directory).

The standard RC files perform first 4.2 mounts, then `nfs` mounts, during booting. On Sun386i systems, `lo` (loopback) mounts are performed just after 4.2 mounts. `/etc/fstab` files depending on alternate mount orders at boot time will fail to work as expected. Manual modification of `/etc/rc.local` will be needed to make such mount orders work.

**WARNINGS**

Loopback mounts must be used with care; the potential for confusing users and applications is enormous. A loopback mount entry in `/etc/fstab` must be placed after the mount points of both directories it depends on. This is most easily accomplished by making the loopback mount entry the last in `/etc/fstab`, though see `mount(8)` for further warnings.

**SEE ALSO**

`chroot(2)`, `mount(2V)`, `fstab(5)`, `mount(8)`

**BUGS**

Because only directories can be mounted or mounted on, the structure of a virtual file system can only be modified at directories.



## NAME

mcp, alm – Sun MCP Multiprotocol Communications Processor/ALM-2 Asynchronous Line Multiplexer

## CONFIG — SUN-3, SUN-4 SYSTEMS

## MCP

```
device mcp0 at vme32d32 ? csr 0x100000 flags 0x1ffff priority 4 vector mcpintr 0x8b
device mcp1 at vme32d32 ? csr 0x1010000 flags 0x1ffff priority 4 vector mcpintr 0x8a
device mcp2 at vme32d32 ? csr 0x1020000 flags 0x1ffff priority 4 vector mcpintr 0x89
device mcp3 at vme32d32 ? csr 0x1030000 flags 0x1ffff priority 4 vector mcpintr 0x88
```

## ALM-2

pseudo-device mcpa64

## CONFIG — SUN-3x SYSTEMS

## MCP

```
device mcp0 at vme32d32 ? csr 0x100000 flags 0x1ffff priority 4 vector mcpintr 0x8b
device mcp1 at vme32d32 ? csr 0x1010000 flags 0x1ffff priority 4 vector mcpintr 0x8a
device mcp2 at vme32d32 ? csr 0x1020000 flags 0x1ffff priority 4 vector mcpintr 0x89
device mcp3 at vme32d32 ? csr 0x1030000 flags 0x1ffff priority 4 vector mcpintr 0x88
device mcp4 at vme32d32 ? csr 0x1040000 flags 0x1ffff priority 4 vector mcpintr 0xa0
device mcp5 at vme32d32 ? csr 0x1050000 flags 0x1ffff priority 4 vector mcpintr 0xa1
device mcp6 at vme32d32 ? csr 0x1060000 flags 0x1ffff priority 4 vector mcpintr 0xa2
device mcp7 at vme32d32 ? csr 0x1070000 flags 0x1ffff priority 4 vector mcpintr 0xa3
```

## ALM-2

pseudo-device mcpa64

## SYNOPSIS

```
#include <fcntl.h>
#include <sys/termios.h>
open("/dev/ttyxy", mode);
open("/dev/ttydn", mode);
open("/dev/cuan", mode);
```

## DESCRIPTION (MCP)

The Sun MCP (Multiprotocol Communications Processor) supports up to four synchronous serial lines in conjunction with SunLink™ Multiple Communication Protocol products.

## DESCRIPTION (ALM-2)

The Sun ALM-2 Asynchronous Line Multiplexer provides 16 asynchronous serial communication lines with modem control and one Centronics-compatible parallel printer port.

Each port supports those `termio(4)` device control functions specified by flags in the `c_cflag` word of the `termios` structure and by the `IGNBRK`, `IGNPAR`, `PARMRK`, or `INPCK` flags in the `c_iflag` word of the `termios` structure are performed by the `mcp` driver. All other `termio(4)` functions must be performed by STREAMS modules pushed atop the driver; when a device is opened, the `ldterm(4M)` and `ttcompat(4M)` STREAMS modules are automatically pushed on top of the stream, providing the standard `termio(4)` interface.

Bit *i* of `flags` may be specified to say that a line is not properly connected, and that the line *i* should be treated as hard-wired with carrier always present. Thus specifying `flags 0x0004` in the specification of `mcp0` would treat line `/dev/ttyh2` in this way.

Minor device numbers in the range 0 – 63 correspond directly to the normal tty lines and are named `/dev/ttyXY`, where *X* represents the physical board as one of the characters `h`, `i`, `j`, or `k`, and *Y* is the line number on the board as a single hexadecimal digit. (Thus the first line on the first board is `/dev/ttyh0`, and the sixteenth line on the third board is `/dev/ttyjf`.)

To allow a single tty line to be connected to a modem and used for both incoming and outgoing calls, a special feature, controlled by the minor device number, has been added. Minor device numbers in the range 128 – 191 correspond to the same physical lines as those above (that is, the same line as the minor device number minus 128).

A dial-in line has a minor device in the range 0 – 63 and is conventionally renamed `/dev/ttydn`, where *n* is a number indicating which dial-in line it is (so that `/dev/ttyd0` is the first dial-in line), and the dial-out line corresponding to that dial-in line has a minor device number 128 greater than the minor device number of the dial-in line and is conventionally named `/dev/cuan`, where *n* is the number of the dial-in line.

The `/dev/cuan` lines are special in that they can be opened even when there is no carrier on the line. Once a `/dev/cuan` line is opened, the corresponding tty line cannot be opened until the `/dev/cuan` line is closed; a blocking open will wait until the `/dev/cuan` line is closed (which will drop Data Terminal Ready, after which Carrier Detect will usually drop as well) and carrier is detected again, and a non-blocking open will return an error. Also, if the `/dev/ttydn` line has been opened successfully (usually only when carrier is recognized on the modem) the corresponding `/dev/cuan` line cannot be opened. This allows a modem to be attached to e.g. `/dev/ttyd0` (renamed from `/dev/ttyh0`) and used for dialin (by enabling the line for login in `/etc/ttytab`) and also used for dialout (by `tip(1C)` or `uucp(1C)`) as `/dev/cua0` when no one is logged in on the line. Note: the bit in the `flags` word in the configuration file (see above) must be zero for this line, which enables hardware carrier detection.

#### IOCTLS

The standard set of `termio ioctl()` calls are supported by the ALM-2.

If the `CRTSCTS` flag in the `c_cflag` is set, output will be generated only if CTS is high; if CTS is low, output will be frozen. If the `CRTSCTS` flag is clear, the state of CTS has no effect. Breaks can be generated by the `TCSBRK`, `TIOCSBRK`, and `TIOCCBRK ioctl()` calls. The modem control lines `TIOCM_CAR`, `TIOCM_CTS`, `TIOCM_RTS`, and `TIOCM_DTR` are provided.

The input and output line speeds may be set to any of the speeds supported by `termio`. The speeds cannot be set independently; when the output speed is set, the input speed is set to the same speed.

#### ERRORS

An `open()` on a `/dev/tty*` or a `/dev/cu*` device will fail if:

- |       |   |
|-------|---|
| ENXIO | The unit being opened does not exist.   |
| EBUSY | The dial-out device is being opened and the dial-in device is already open, or the dial-in device is being opened with a no-delay open and the dial-out device is already open. |
| EBUSY | The unit has been marked as exclusive-use by another process with a <code>TIOCEXCL ioctl()</code> call.   |
| EINTR | The open was interrupted by the delivery of a signal.   |

#### DESCRIPTION (PRINTER PORT)

The printer port is Centronics-compatible and is suitable for most common parallel printers. Devices attached to this interface are normally handled by the line printer spooling system, and should not be accessed directly by the user.

Minor device numbers in the range 64 – 67 access the printer port, and the recommended naming is `/dev/mcpp[0-3]`.

#### IOCTLS

Various control flags and status bits may be fetched and set on an MCP printer port. The following flags and status bits are supported; they are defined in `sundev/mcpcmd.h`:

MCPRIGNSLCT	0x02	set if interface ignoring SLCT— on open
MCPRDIAG	0x04	set if printer is in self-test mode
MCPRVMEINT	0x08	set if VME bus interrupts enabled
MCPRIPTPE	0x10	print message when out of paper
MCPRIPTSLCT	0x20	print message when printer offline

MCPRPE	0x40	set if device ready, cleared if device out of paper
MCPRSLCT	0x80	set if device online (Centronics SLCT asserted)

The flags **MCPRINTSLCT**, **MCPRINTPE**, and **MCPRDIAG** may be changed; the other bits are status bits and may not be changed.

The **ioctl()** calls supported by MCP printer ports are listed below.

MCPIOGPR	The argument is a pointer to an <b>unsigned char</b> . The printer flags and status bits are stored in the <b>unsigned char</b> pointed to by the argument.
MCPIOSPR	The argument is a pointer to an <b>unsigned char</b> . The printer flags are set from the <b>unsigned char</b> pointed to by the argument.

#### ERRORS

Normally, the interface only reports the status of the device when attempting an **open(2V)** call. An **open()** on a **/dev/mcpp\*** device will fail if:

ENXIO	The unit being opened does not exist.
EIO	The device is offline or out of paper.

Bit 17 of the configuration flags may be specified to say that the interface should ignore Centronics SLCT- and RDY/PE- when attempting to open the device, but this is normally useful only for configuration and troubleshooting; if the SLCT- and RDY lines are not asserted during an actual data transfer (as with a **write(2V)** call), no data is transferred.

#### FILES

<b>/dev/mcpp[0-3]</b>	parallel printer port
<b>/dev/tty[h-k][0-9a-f]</b>	hardwired tty lines
<b>/dev/ttyd[0-9a-f]</b>	dialin tty lines
<b>/dev/cua[0-9a-f]</b>	dialout tty lines

#### SEE ALSO

**tip(1C)**, **uucp(1C)**, **mti(4S)**, **termio(4)**, **ldterm(4M)**, **ttcompat(4M)**, **zs(4S)**, **ttysoftcar(8)**

#### DIAGNOSTICS

Most of these diagnostics "should never happen;" their occurrence usually indicates problems elsewhere in the system as well.

##### **mcpn: silo overflow.**

More than *n* characters (*n* very large) have been received by the **mcp** hardware without being read by the software.

##### **\*\*\*port n supports RS449 interface\*\*\***

Probably an incorrect jumper configuration. Consult the hardware manual.

##### **mcp port n receive buffer error**

The **mcp** encountered an error concerning the synchronous receive buffer.

##### **Printer on mcppn is out of paper**

##### **Printer on mcppn paper ok**

##### **Printer on mcppn is offline**

##### **Printer on mcppn online**

Assorted printer diagnostics, if enabled as discussed above.

#### BUGS

Note: pin 4 is used for hardware flow control on ALM-2 ports 0 through 3. These two pins should *not* be tied together on the ALM end.

**NAME**

mem, kmem, zero, vme16d16, vme24d16, vme32d16, vme16d32, vme24d32, vme32d32, eeprom, atbus, sbus – main memory and bus I/O space

**CONFIG**

None; included with standard system.

**DESCRIPTION**

These devices are special files that map memory and bus I/O space. They may be read, written, seeked and (except for **kmem**) memory-mapped. See **read(2V)**, **write(2V)**, **mmap(2)**, and **directory(3V)**.

**All Systems**

**mem** is a special file that is an image of the physical memory of the computer. It may be used, for example, to examine (and even to patch) the system.

**kmem** is a special file that is an image of the kernel virtual memory of the system.

**zero** is a special file which is a source of private zero pages.

**eeprom** is a special file that is an image of the EEPROM or NVRAM.

**Sun-3 and Sun-4 Systems VMEbus**

**vme16d16** (also known as **vme16**) is a special file that is an image of VMEbus 16-bit addresses with 16-bit data. **vme16** address space extends from 0 to 64K.

**vme24d16** (also known as **vme24**) is a special file that is an image of VMEbus 24-bit addresses with 16-bit data. **vme24** address space extends from 0 to 16 Megabytes. The VME 16-bit address space overlaps the top 64K of the 24-bit address space.

**vme32d16** is a special file that is an image of VMEbus 32-bit addresses with 16-bit data.

**vme16d32** is a special file that is an image of VMEbus 16-bit addresses with 32-bit data.

**vme24d32** is a special file that is an image of VMEbus 24-bit addresses with 32-bit data.

**vme32d32** (also known as **vme32**) is a special file that is an image of VMEbus 32-bit addresses with 32-bit data. **vme32** address space extends from 0 to 4 Gigabytes. The VME 24-bit address space overlaps the top 16 Megabytes of the 32-bit address space.

**SPARCstation 1 Systems**

The **sbus** is represented by a series of entries each of which is an image of a single **sbus** slot. The entries are named **sbus $n$** , where  $n$  is the slot number in hexadecimal. The number of **sbus** slots and the address range within each slot may vary between implementations.

**Sun386i Systems**

**atbus** is a special file that is an image of the AT bus space. It extends from 0 to 16 Megabytes.

**FILES**

**/dev/mem**  
**/dev/kmem**  
**/dev/zero**  
**/dev/vme16d16**  
**/dev/vme16**  
**/dev/vme24d16**  
**/dev/vme24**  
**/dev/vme32d16**  
**/dev/vme16d32**  
**/dev/vme24d32**  
**/dev/vme32d32**  
**/dev/vme32**  
**/dev/eeprom**  
**/dev/atbus**  
**/dev/sbus[0-3]**

**SEE ALSO**

**mmap(2), read(2V), write(2V), directory(3V)**

**NAME**

mouse – Sun mouse

**CONFIG**

None; included in standard system.

**DESCRIPTION**

The **mouse** indirect device provides access to the Sun Workstation mouse. When opened, it redirects operations to the standard mouse device for the workstation (attached either to a CPU serial or parallel port), and pushes the **ms(4M)** and **ttcompat(4M)** STREAMS modules on top of that device.

**FILES**

**/dev/mouse**

**SEE ALSO**

**ms(4M)**, **ttcompat(4M)**, **win(4S)**, **zs(4S)**

**NAME**

`ms` – Sun mouse STREAMS module

**CONFIG**

`pseudo-device`*msn*

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/time.h>
#include <sys/stream.h>
#include <sys/stropts.h>
#include <sundev/vuid_event.h>
#include <sundev/msio.h>
ioctl(fd, I_PUSH, "ms");
```

**DESCRIPTION**

The `ms` STREAMS module processes byte streams generated by mice attached to a CPU serial or parallel port. When this module is pushed onto a stream, it sends a `TCSETS` `ioctl` downstream, setting the baud rate to 1200 baud and the character size to 8 bits, and enabling the receiver. All other flag words are cleared. It assumes only that the `termios(3V)` functions provided by the `zs(4S)` driver are supported; no other functions need be supported.

The mouse is expected to generate a stream of bytes encoding mouse motions and changes in the state of the buttons.

Each mouse sample in the byte stream consists of three bytes: the first byte gives the button state with value `0x87|but`, where `but` is the low three bits giving the mouse buttons, where a 0 (zero) bit means that a button is pressed, and a 1 (one) bit means a button is not pressed. Thus if the left button is down the value of this sample is `0x83`, while if the right button is down the byte is `0x86`.

The next two bytes of each sample give the `x` and `y` deltas of this sample as signed bytes. The mouse uses a lower-left coordinate system, so moves to the right on the screen yield positive `x` values and moves down the screen yield negative `y` values.

The beginning of a sample is identifiable because the delta's are constrained to not have values in the range `0x80-0x87`.

A stream with `ms` pushed onto it can be used as a device that emits `firm_events` as specified by the protocol of a *Virtual User Input Device*. It understands `VUIDSFORMAT`, `VUIDGFORMAT`, `VUIDSADDR` and `VUIDGADDR` `ioctls` (see reference below).

**IOCTLS**

`ms` responds to the following `ioctls`, as defined in `<sundev/msio.h>` and `<sundev/vuid_event.h>`. All other `ioctls` are passed downstream. As `ms` sets the parameters of the serial port when it is opened, no `termios(3V)` `ioctls` should be performed on a stream with `ms` on it, as `ms` expects the device parameters to remain as it set them.

The `MSIOGETPARMS` and `MSIOSETPARMS` calls use a structure of type `Ms_parms`, which is a structure defined in `<sundev/msio.h>`:

```
typedef struct {
    int    jitter_thresh;
    int    speed_law;
    int    speed_limit;
} Ms_parms;
```

*jitter\_thresh* is the “jitter threshold” of the mouse. Motions of fewer than *jitter\_thresh* units along both axes that occur in less than 1/12 second are treated as “jitter” and ignored. Thus, if the mouse moves fewer than *jitter\_thresh* units and then moves back to its original position in less than 1/12 of a second, the motion is considered to be “noise” and ignored. If it moves fewer than *jitter\_thresh* units and continues to move so that it has not returned to its original position after 1/12 of a second, the motion is considered to be real and is reported.

*speed\_limit* indicates whether extremely large motions are to be ignored. If it is 1, a “speed limit” is applied to mouse motions; motions along either axis of more than *speed\_limit* units are discarded.

Note: these parameters are global; if they are set for any mouse on a workstation, they apply to any other mice attached to that workstation as well.

#### **VIDSFORMAT**

#### **VIDGFORMAT**

#### **VIDSADDR**

#### **VIDGADDR**

These are standard *Virtual User Input Device ioctls*. See *SunView System Programmer's Guide* for a description of their operation.

#### **MSIOGETPARMS**

The argument is a pointer to a *Ms\_parms*. The current mouse parameters are stored in that structure.

#### **MSIOSETPARMS**

The argument is a pointer to a *ms\_parms*. The current mouse parameters are set from the values in that structure.

#### **SEE ALSO**

*mouse(4S)*, *termios(3V)*, *win(4S)*, *zs(4S)*

*SunView System Programmer's Guide*



**NAME**

mti – Systech MTI-800/1600 multi-terminal interface

**CONFIG — SUN-3, SUN-3x, SUN-4 SYSTEMS**

```
device mti0 at vme16d16 ? csr 0x620 flags 0xffff priority 4 vector mtiintr 0x88
device mti1 at vme16d16 ? csr 0x640 flags 0xffff priority 4 vector mtiintr 0x89
device mti2 at vme16d16 ? csr 0x660 flags 0xffff priority 4 vector mtiintr 0x8a
device mti3 at vme16d16 ? csr 0x680 flags 0xffff priority 4 vector mtiintr 0x8b
```

**SYNOPSIS**

```
#include <fcntl.h>
#include <sys/termios.h>
open("/dev/ttyxy", mode);
open("/dev/ttydn", mode);
open("/dev/cuan", mode);
```

**DESCRIPTION**

The Systech MTI card provides 8 (MTI-800) or 16 (MTI-1600) serial communication lines with modem control. Each port supports those `termio(4)` device control functions specified by flags in the `c_cflag` word of the `termios` structure and by the `IGNBRK`, `IGNPAR`, `PARMRK`, or `INPCK` flags in the `c_iflag` word of the `termios` structure are performed by the `mti` driver. All other `termio(4)` functions must be performed by STREAMS modules pushed on top of the driver; when a device is opened, the `ldterm(4M)` and `ttcompat(4M)` STREAMS modules are automatically pushed on top of the stream, providing the standard `termio(4)` interface.

Bit *i* of `flags` may be specified to say that a line is not properly connected, and that the line *i* should be treated as hard-wired with carrier always present. Thus specifying `flags 0x0004` in the specification of `mti0` would treat line `/dev/tty02` in this way.

Minor device numbers in the range 0 – 63 correspond directly to the normal tty lines and are named `/dev/ttyXY`, where *X* is the physical board number (0 – 3), and *Y* is the line number on the board as a single hexadecimal digit. Thus the first line on the first board is `/dev/tty00`, and the sixteenth line on the third board is `/dev/tty2f`.

To allow a single tty line to be connected to a modem and used for both incoming and outgoing calls, a special feature, controlled by the minor device number, has been added. Minor device numbers in the range 128 – 191 correspond to the same physical lines as those above (that is, the same line as the minor device number minus 128).

A dial-in line has a minor device in the range 0 – 63 and is conventionally renamed `/dev/ttydn`, where *n* is a number indicating which dial-in line it is (so that `/dev/ttyd0` is the first dial-in line), and the dial-out line corresponding to that dial-in line has a minor device number 128 greater than the minor device number of the dial-in line and is conventionally named `/dev/cuan`, where *n* is the number of the dial-in line.

The `/dev/cuan` lines are special in that they can be opened even when there is no carrier on the line. Once a `/dev/cuan` line is opened, the corresponding tty line can not be opened until the `/dev/cuan` line is closed; a blocking open will wait until the `/dev/cuan` line is closed (which will drop Data Terminal Ready, after which Carrier Detect will usually drop as well) and carrier is detected again, and a non-blocking open will return an error. Also, if the `/dev/ttydn` line has been opened successfully (usually only when carrier is recognized on the modem) the corresponding `/dev/cuan` line can not be opened. This allows a modem to be attached to for example, `/dev/ttyd0` (renamed from `/dev/tty00`) and used for dial-in (by enabling the line for login in `/etc/ttytab`) and also used for dial-out (by `tip(1C)` or `uucp(1C)`) as `/dev/cua0` when no one is logged in on the line. Note: the bit in the `flags` word in the configuration file (see above) must be zero for this line, which enables hardware carrier detection.

**WIRING**

The Systech requires the CTS modem control signal to operate. If the device does not supply CTS then RTS should be jumpered to CTS at the distribution panel (short pins 4 to 5). Also, the CD (carrier detect) line does not work properly. When connecting a modem, the modem's CD line should be wired to DSR, which the software will treat as carrier detect.

**IOCTLS**

The standard set of `termio ioctl()` calls are supported by `mti`.

The state of the `CRTSCTS` flag in the `c_cflag` word has no effect; no output will be generated unless CTS is high. Breaks can be generated by the `TCSBRK`, `TIOCSBRK`, and `TIOCCBRK ioctl()` calls. The modem control lines `TIOCM_CAR`, `TIOCM_CTS`, `TIOCM_RTS`, and `TIOCM_DTR` are provided; however, as described above, the DSR line is treated as CD and the CD line is ignored.

The input and output line speeds may be set to any of the speeds supported by `termio`. The speeds cannot be set independently; when the output speed is set, the input speed is set to the same speed. The baud rates `B200` and `B38400` are not supported by the hardware; `B200` selects 2000 baud, and `B38400` selects 7200 baud.

**ERRORS**

An `open()` will fail if:

<code>ENXIO</code>	The unit being opened does not exist.
<code>EBUSY</code>	The dial-out device is being opened and the dial-in device is already open, or the dial-in device is being opened with a no-delay open and the dial-out device is already open.
<code>EBUSY</code>	The unit has been marked as exclusive-use by another process with a <code>TIOCEXCL ioctl()</code> call.
<code>EINTR</code>	The open was interrupted by the delivery of a signal.

**FILES**

<code>/dev/tty[0-3][0-9a-f]</code>	hardwired tty lines
<code>/dev/ttyd[0-9a-f]</code>	dial-in tty lines
<code>/dev/cua[0-9a-f]</code>	dial-out tty lines

**SEE ALSO**

`tip(1C)`, `uucp(1C)`, `mcp(4S)`, `termio(4)`, `ldterm(4M)`, `ttcompat(4M)`, `zs(4S)`, `ttysoftcar(8)`

**DIAGNOSTICS**

Most of these diagnostics "should never happen" and their occurrence usually indicates problems elsewhere in the system.

**`mtin, n`: silo overflow.**

More than 512 characters have been received by the `mti` hardware without being read by the software. Extremely unlikely to occur.

**`mtin`: read error code `<n>`. Probable hardware fault**

The `mti` returned the indicated error code. See the MTI manual.

**`mtin`: DMA output error.**

The `mti` encountered an error while trying to do DMA output.

**`mtin`: impossible response `n`.**

The `mti` returned an error it could not understand.

**NAME**

mtio – general magnetic tape interface

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/ioctl.h>
#include <sys/mtio.h>
```

**DESCRIPTION**

1/2", 1/4" and 8 mm magnetic tape drives all share the same general character device interface.

There are two types of tape records: data records and end-of-file (EOF) records. EOF records are also known as tape marks and file marks. A record is separated by interrecord (or tape) gaps on a tape.

End-of-recorded-media (EOM) is indicated by two EOF marks on 1/2" tape; by one on 1/4" and 8 mm cartridge tapes.

**1/2" Reel Tape**

Data bytes are recorded in parallel onto the 9-track tape. The number of bytes in a physical record varies between 1 and 65535 bytes.

The recording formats available (check specific tape drive) are 800 BPI, 1600 BPI, and 6250 BPI, and data compression. Actual storage capacity is a function of the recording format and the length of the tape reel. For example, using a 2400 foot tape, 20 MB can be stored using 800 BPI, 40 MB using 1600 BPI, 140 MB using 6250 BPI, or up to 700 MB using data compression.

**1/4" Cartridge Tape**

Data is recorded serially onto 1/4" cartridge tape. The number of bytes per record is determined by the physical record size of the device. The I/O request size must be a multiple of the physical record size of the device. For QIC-11, QIC-24, and QIC-150 tape drives the block size is 512 bytes.

The records are recorded on tracks in a serpentine motion. As one track is completed, the drive switches to the next and begins writing in the opposite direction, eliminating the wasted motion of rewinding. Each file, including the last, ends with one file mark.

Storage capacity is based on the number of tracks the drive is capable of recording. For example, 4-track drives can only record 20 MB of data on a 450 foot tape; 9-track drives can record up to 45 MB of data on a tape of the same length. QIC-11 is the only tape format available for 4-track tape drives. In contrast, 9-track tape drives can use either QIC-24 or QIC-11. Storage capacity is not appreciably affected by using either format. QIC-24 is preferable to QIC-11 because it records a reference signal to mark the position of the first track on the tape, and each block has a unique block number.

The QIC-150 tape drives require DC-6150 (or equivalent) tape cartridges for writing. However, they can read other tape cartridges in QIC-11, QIC-24, QIC-120, or QIC-150 tape formats.

**8 mm Cartridge Tape**

Data is recorded serially onto 8 mm helical scan cartridge tape. The number of bytes in a physical record varies between 1 and 65535 bytes. Currently one density is available.

**Read Operation**

`read(2V)` reads the next record on the tape. The record size is passed back as the number of bytes read, provided it is no greater than the number requested. When a tape mark is read, a zero byte count is returned; another read will fetch the first record of the next tape file. Two successive reads returning zero byte counts indicate the EOM. No further reading should be performed past the EOM.

Fixed-length I/O tape devices require the number of bytes read to be a multiple of the physical record size. For example, 1/4" cartridge tape devices only read multiples of 512 bytes. If the blocking factor is greater than 64512 bytes (minphys limit), fixed-length I/O tape devices read multiple records.

Tape devices which support variable-length I/O operations, such as 1/2" and 8mm tape, may read a range of 1 to 65535 bytes. If the record size exceeds 65535 bytes, the driver reads multiple records to satisfy the request. These multiple records are limited to 65534 bytes.

**Write Operation**

**write(2V)** writes the next record on the tape. The record has the same length as the given buffer.

Writing is allowed on 1/4" tape at either the beginning of tape or after the last written file on the tape.

Writing is not so restricted on 1/2" and 8 mm cartridge tape. Care should be used when appending files onto 1/2" reel tape devices, since an extra file mark is appended after the last file to mark the EOM. This extra file mark must be overwritten to prevent the creation of a null file. To facilitate write append operations, a space to the EOM ioctl is provided. Care should be taken when overwriting records; the erase head is just forward of the write head and any following records will also be erased.

Fixed-length I/O tape devices require the number of bytes written to be a multiple of the physical record size. For example, 1/4" cartridge tape devices only write multiples of 512 bytes. Fixed-length I/O tape devices write multiple records if the blocking factor is greater than 64512 bytes (minphys limit). These multiple writes are limited to 64512 bytes. For example, if a write request is issued for 65536 bytes using a 1/4" cartridge tape, two writes are issued; the first for 64512 bytes and the second for 1024 bytes.

Tape devices which support variable-length I/O operations, such as 1/2" and 8mm tape, may write a range of 1 to 65535 bytes. If the record size exceeds 65535 bytes, the driver writes multiple records to satisfy the request. These multiple records are limited to 65534 bytes. As an example, if a write request for 65540 bytes is issued using 1/2" reel tape, two records are written; one for 65534 bytes followed by one for 6 bytes.

EOT handling on write is different among the various devices; see the appropriate device manual page. Reading past EOT is transparent to the user.

Seeks are ignored in tape I/O.

**Close Operation**

Magnetic tapes are rewound when closed, except when the "no-rewind" devices have been specified. The names of no-rewind device files use the letter **n** as the beginning of the final component. The no-rewind version of `/dev/rmt0` is `/dev/nrmt0`.

If data was written, a file mark is automatically written by the driver upon close. If the rewinding device was specified, the tape will be rewound after the file mark is written. If the user wrote a file mark prior to closing, then no file mark is written upon close. If a file positioning ioctl, like `rewind`, is issued after writing, a file mark is written before repositioning the tape.

Note: for 1/2" reel tape devices, two file marks are written to mark the EOM before rewinding or performing a file positioning ioctl. If the user wrote a file mark before closing a 1/2" reel tape device, the driver will always write a file mark before closing to insure that the end of recorded media is marked properly. If the non-rewinding `xt` device was specified, two file marks are written and the tape is left positioned between the two so that the second one is overwritten on a subsequent `open(2V)` and `write(2V)`. For performance reasons, the `st` driver postpones writing the second tape mark until just before a file positioning ioctl is issued (for example, `rewind`). This means that the user must not manually rewind the tape because the tape will be missing the second tape mark which marks EOM.

If no data was written and the driver was opened for WRITE-ONLY access, a file mark is written thus creating a null file.

**Ioctls**

Not all devices support all ioctls. The driver returns an ENOTTY error on unsupported ioctls.

The following structure definitions for magnetic tape ioctl commands are from `<sys/mtio.h>`:

```
/* structure for MTIOCTOP – magnetic tape operation command */
struct mtop {
    short mt_op;          /* operation */
    daddr_t mt_count;    /* number of operations */
};
```

The following ioctls are supported:

MTWEOF	write an end-of-file record
MTFSF	forward space over file mark
MTBSF	backward space over file mark (1/2", 8 mm only)
MTFSR	forward space to inter-record gap
MTBSR	backward space to inter-record gap
MTREW	rewind
MTOFFL	rewind and take the drive offline
MTNOP	no operation, sets status only
MTRETEN	retension the tape (cartridge tape only)
MTERASE	erase the entire tape and rewind
MTEOM	position to EOM
MTNBSF	backward space file to beginning of file

```

/* structure for MTIOCGET – magnetic tape get status command */
struct mtget {
    short    mt_type;           /* type of magtape device */

    /* the following two registers are device dependent */
    short    mt_dsreg;         /* "drive status" register */
    short    mt_erreg;         /* "error" register */

    /* optional error info. */
    daddr_t  mt_resid;         /* residual count */
    daddr_t  mt_fileno;        /* file number of current position */
    daddr_t  mt_blkno;         /* block number of current position */
    u_short  mt_flags;
    short    mt_bf;           /* optimum blocking factor */
};

```

When spacing forward over a record (either data or EOF), the tape head is positioned in the tape gap between the record just skipped and the next record. When spacing forward over file marks (EOF records), the tape head is positioned in the tape gap between the next EOF record and the record that follows it.

When spacing backward over a record (either data or EOF), the tape head is positioned in the tape gap immediately preceding the tape record where the tape head is currently positioned. When spacing backward over file marks (EOF records), the tape head is positioned in the tape gap preceding the EOF. Thus the next read would fetch the EOF.

Note, the following features are unique to the *st* driver: record skipping does not go past a file mark; file skipping does not go past the EOM. Both the *st* and *xt* drivers stop upon encountering EOF during a record skipping command, but leave the tape positioned differently. For example, after an *MTFSR* <huge number> command the *st* driver leaves the tape positioned *before* the EOF. After the same command, the *xt* driver leaves the tapes positioned *after* the EOF. Consequently on the next read, the *xt* driver fetches the first record of the next file whereas the *st* driver fetches the EOF. A related *st* feature is that EOFs remain pending until the tape is closed. For example, a program which first reads all the records of a file up to and including the EOF and then performs an *MTFSF* command will leave the tape positioned just after that same EOF, rather than skipping the next file.

The *MTNBSF* and *MTFSF* operations are inverses. Thus, an *MTFSF* “-1” is equivalent to an *MTNBSF* “1”. An *MTNBSF* “0” is the same as *MTFSF* “0”; both position the tape device to the beginning of the current file.

*MTBSF* moves the tape backwards by file marks. The tape position will end on the beginning of tape side of the desired file mark.

MTBSR and MTFSR operations perform much like space file operations, except that they move by records instead of files. Variable-length I/O devices (1/2" reel, for example) space actual records; fixed-length I/O devices space physical records (blocks). 1/4" cartridge tape, for example, spaces 512 byte physical records. The status ioctl residual count contains the number of files or records not skipped.

MTOFFL rewinds and, if appropriate, takes the device offline by unloading the tape. The tape must be inserted before the tape device can be used again.

MTRETEN The retension ioctl only applies to 1/4" cartridge tape devices. It is used to restore tape tension improving the tape's soft error rate after extensive start-stop operations or long-term storage.

MTERASE rewinds the tape, erases it completely, and returns to the beginning of tape.

MTEOM positions the tape at a location just after the last file written on the tape. For 1/4" cartridge and 8 mm tape, this is after the last file mark on the tape. For 1/2" reel tape, this is just after the first file mark but before the second (and last) file mark on the tape. Additional files can then be appended onto the tape from that point.

Note the difference between MTBSF (backspace over file mark) and MTNBSF (backspace file to beginning of file). The former moves the tape backward until it crosses an EOF mark, leaving the tape positioned *before* the file mark. The latter leaves the tape positioned *after* the file mark. Hence, "MTNBSF n" is equivalent to "MTBSF (n+1)" followed by "MTFSF 1". 1/4" cartridge tape devices do not support MTBSF.

The MTIOCGET get status ioctl call returns the drive id (*mt\_type*), sense key error (*mt\_erreg*), file number (*mt\_fileno*), optimum blocking factor (*mt\_bf*) and record number (*mt\_blkno*) of the last error. The residual count (*mt\_resid*) is set to the number of bytes not transferred or files/records not spaced. The flags word (*mt\_flags*) contains information such as whether the device is SCSI, whether it is a reel device and whether the device supports absolute file positioning.

#### EXAMPLES

Suppose you have written 3 files to the non-rewinding 1/2" tape device, `/dev/nrmt0`, and that you want to go back and `dd(1)` the second file off the tape. The commands to do this are:

```
mt -f /dev/nrmt0 bsf 3
mt -f /dev/nrmt0 fsf 1
dd if=/dev/nrmt0
```

To accomplish the same tape positioning in a C program, followed by a get status ioctl:

```
struct mtop mt_command;
struct mtget mt_status;

mt_command.mt_op = MTBSF;
mt_command.mt_count = 3;
ioctl(fd, MTIOCTOP, &mt_command);
mt_command.mt_op = MTFSF;
mt_command.mt_count = 1;
ioctl(fd, MTIOCTOP, &mt_command);
ioctl(fd, MTIOCGET, (char *)&mt_status);
```

or

```
struct mtop mt_command;
struct mtget mt_status;

mt_command.mt_op = MTNBSF;
mt_command.mt_count = 2;
ioctl(fd, MTIOCTOP, &mt_command);
ioctl(fd, MTIOCGET, (char *)&mt_status);
```

**FILES**

**/dev/rmt\***  
**/dev/rst\***  
**/dev/rar\***  
**/dev/nrmt\***  
**/dev/nrst\***  
**/dev/nrar\***

**SEE ALSO**

**dd(1), mt(1), tar(1), read(2V), write(2V), ar(4S), st(4S), tm(4S), xt(4S)**

*1/4 Inch Tape Drive Tutorial*

**WARNINGS**

Avoid the use of device files **/dev/rmt4** and **/dev/rmt12**, as they are going away in a future release.

**NAME**

nfs, NFS – network file system

**CONFIG**

options NFS

**DESCRIPTION**

The Network File System, or NFS, allows a client workstation to perform transparent file access over the network. Using it, a client workstation can operate on files that reside on a variety of servers, server architectures and across a variety of operating systems. Client file access calls are converted to NFS protocol requests, and are sent to the server system over the network. The server receives the request, performs the actual file system operation, and sends a response back to the client.

The Network File System operates in a stateless fashion using remote procedure (RPC) calls built on top of external data representation (XDR) protocol. These protocols are documented in *Network Programming*. The RPC protocol provides for version and authentication parameters to be exchanged for security over the network.

A server can grant access to a specific filesystem to certain clients by adding an entry for that filesystem to the server's `/etc/exports` file and running `exportfs(8)`.

A client gains access to that filesystem with the `mount(2V)` system call, which requests a file handle for the filesystem itself. Once the filesystem is mounted by the client, the server issues a file handle to the client for each file (or directory) the client accesses or creates. If the file is somehow removed on the server side, the file handle becomes stale (dissociated with a known file).

A server may also be a client with respect to filesystems it has mounted over the network, but its clients cannot gain access to those filesystems. Instead, the client must mount a filesystem directly from the server on which it resides.

The user ID and group ID mappings must be the same between client and server. However, the server maps uid 0 (the super-user) to uid -2 before performing access checks for a client. This inhibits super-user privileges on remote filesystems. This may be changed by use of the “anon” export option. See `exportfs(8)`.

NFS-related routines and structure definitions are described in *Network Programming*.

**ERRORS**

Generally physical disk I/O errors detected at the server are returned to the client for action. If the server is down or inaccessible, the client will see the console message:

**NFS server *host* not responding still trying.**

Depending on whether the file system has been mounted “hard” or “soft” (see `mount(8)`), the client will either continue (forever) to resend the request until it receives an acknowledgement from the server, or return an error to user-level. For hard mounts, this means the server can crash or power down and come back up without any special action required by the client. If the “intr” mount option was not specified, a client process requesting I/O will block and remain insensitive to signals, sleeping inside the kernel at `PRI-BIO` until the request is satisfied.

**FILES**

`/etc/exports`

**SEE ALSO**

`mount(2V)`, `exports(5)`, `fstab(5)`, `fstab(5)`, `exportfs(8)`, `mount(8)`, `nfsd(8)`, `sticky(8)`

*Network Programming*



**BUGS**

When a file that is opened by a client is unlinked (by the server), a file with a name of the form `.nfsXXX` (where `XXX` is a number) is created by the client. When the open file is closed, the `.nfsXXX` file is removed. If the client crashes before the file can be closed, the `.nfsXXX` file is not removed.

NFS servers usually mark their clients' swap files specially to avoid being required to sync their inodes to disk before returning from writes. See `sticky(8)`.

**NAME**

nit – Network Interface Tap

**CONFIG**

```

pseudo-device  clone
pseudo-device  snit
pseudo-device  pf
pseudo-device  nbuf

```

**SYNOPSIS**

```

#include <sys/file.h>
#include <sys/ioctl.h>
#include <net/nit_pf.h>
#include <net/nit_buf.h>

fd = open("/dev/nit", mode);
ioctl(fd, I_PUSH, "pf");
ioctl(fd, I_PUSH, "nbuf");

```

**DESCRIPTION**

NIT (the Network Interface Tap) is a facility composed of several STREAMS modules and drivers. These components collectively provide facilities for constructing applications that require link-level network access. Examples of such applications include **rarpd**(8C), which is a user-level implementation of the Reverse ARP protocol, and **etherfind**(8C), which is a network monitoring and trouble-shooting program.

NIT consists of several components that are summarized below. See their Reference Manual entries for detailed information about their specification and operation.

**nit\_if**(4M) This component is a STREAMS device driver that interacts directly with the system's Ethernet drivers. After opening an instance of this device it must be bound to a specific Ethernet interface before becoming usable. Subsequently, **nit\_if** transcribes packets arriving on the interface to the read side of its associated stream and delivers messages reaching it on the write side of its stream to the raw packet output code for transmission over the interface.

**nit\_pf**(4M) This module provides packet-filtering services, allowing uninteresting incoming packets to be discarded with minimal loss of efficiency. It passes through unaltered all outgoing messages (those on the stream's write side).

**nit\_buf**(4M) This module buffers incoming messages into larger aggregates, thereby reducing the overhead incurred by repeated **read**(2V) system calls.

NIT clients mix and match these components, based on their particular requirements. For example, the reverse ARP daemon concerns itself only with packets of a specific type and deals with low traffic volumes. Thus, it uses **nit\_if** for access to the network and **nit\_pf** to filter out all incoming packets except reverse ARP packets, but omits the **nit\_buf** buffering module since traffic is not high enough to justify the additional complexity of unpacking buffered packets. On the other hand, the **etherd**(8C) program, which collects Ethernet statistics for **traffic**(1C) to display, must examine every packet on the network. Therefore, it omits the **nit\_pf** module, since there is nothing it wishes to screen out, and includes the **nit\_buf** module, since most networks have very heavy aggregate packet traffic.

**EXAMPLES**

The following code fragments outline how to program against parts of the NIT interface. For the sake of brevity, all error-handling code has been elided.

**initdevice** comes from **etherfind** and sets up its input stream configuration.

```

initdevice(if_flags, snaplen, chunksize)
    u_long  if_flags,
           snaplen,
           chunksize;

```

```

{
    struct strioctl    si;
    struct ifreq       ifr;
    struct timeval     timeout;

    if_fd = open(NIT_DEV, O_RDONLY);

    /* Arrange to get discrete messages from the stream. */
    ioctl(if_fd, I_SRDOPT, (char *)RMSGD);

    si.ic_timeout = INFTIM;

    /* Push and configure the buffering module. */
    ioctl(if_fd, I_PUSH, "nbuf");

    timeout.tv_sec = 1;
    timeout.tv_usec = 0;
    si.ic_cmd = NIOCSTIME;
    si.ic_len = sizeof timeout;
    si.ic_dp = (char *)&timeout;
    ioctl(if_fd, I_STR, (char *)&si);

    si.ic_cmd = NIOCSCHUNK;
    si.ic_len = sizeof chunksize;
    si.ic_dp = (char *)&chunksize;
    ioctl(if_fd, I_STR, (char *)&si);

    /* Configure the nit device, binding it to the proper
       underlying interface, setting the snapshot length,
       and setting nit_if-level flags. */
    strncpy(ifr.ifr_name, device, sizeof ifr.ifr_name);
    ifr.ifr_name[sizeof ifr.ifr_name - 1] = '\0';
    si.ic_cmd = NIOCBIND;
    si.ic_len = sizeof ifr;
    si.ic_dp = (char *)&ifr;
    ioctl(if_fd, I_STR, (char *)&si);

    if (snaplen > 0) {
        si.ic_cmd = NIOCSSNAP;
        si.ic_len = sizeof snaplen;
        si.ic_dp = (char *)&snaplen;
        ioctl(if_fd, I_STR, (char *)&si);
    }

    if (if_flags != 0) {
        si.ic_cmd = NIOCSFLAGS;
        si.ic_len = sizeof if_flags;
        si.ic_dp = (char *)&if_flags;
        ioctl(if_fd, I_STR, (char *)&si);
    }

    /* Flush the read queue, to get rid of anything that accumulated
       before the device reached its final configuration. */
    ioctl(if_fd, I_FLUSH, (char *)FLUSHR);
}

```

Here is the skeleton of the packet reading loop from **etherfind**. It illustrates how to cope with dismantling the headers the various NIT components glue on.

```

while ((cc = read(if_fd, buf, chunksize)) >= 0) {
    register u_char    *bp = buf,
                      *bufstop = buf + cc;

    /* Loop through each message in the chunk. */
    while (bp < bufstop) {
        register u_char    *cp = bp;
        struct nit_bufhdr *hdrp;
        struct timeval     *tvp = NULL;
        u_long             drops = 0;
        u_long             pktlen;

        /* Extract information from the successive objects
           embedded in the current message. Which ones we
           have depends on how we set up the stream (and
           therefore on what command line flags were set).

           If snaplen is positive then the packet was truncated
           before the buffering module saw it, so we must
           obtain its length from the nit_if-level nit_iflen
           header. Otherwise the value in *hdrp suffices. */
        hdrp = (struct nit_bufhdr *)cp;
        cp += sizeof *hdrp;
        if (tflag) {
            struct nit_iftime *ntp;

            ntp = (struct nit_iftime *)cp;
            cp += sizeof *ntp;

            tvp = &ntp->nh_timestamp;
        }
        if (dflag) {
            struct nit_ifdrops *ndp;

            ndp = (struct nit_ifdrops *)cp;
            cp += sizeof *ndp;

            drops = ndp->nh_drops;
        }
        if (snaplen > 0) {
            struct nit_iflen *nlp;

            nlp = (struct nit_iflen *)cp;
            cp += sizeof *nlp;

            pktlen = nlp->nh_pktlen;
        }
        else
            pktlen = hdrp->nhb_msglen;

        sp = (struct sample *)cp;
        bp += hdrp->nhb_totlen;

        /* Process the packet. */
    }
}

```

**FILES**

**/dev/nit**                    clone device instance referring to **nit\_if**

**SEE ALSO**

**traffic(1C), read(2V), nit\_if(4M), nit\_pf(4M), nit\_buf(4M), etherd(8C), etherfind(8C), rarpd(8C)**

**NAME**

`nit_buf` – STREAMS NIT buffering module

**CONFIG**

`pseudo-device nbuf`

**SYNOPSIS**

```
#include <sys/ioctl.h>
#include <net/nit_buf.h>
ioctl(fd, I_PUSH, "nbuf");
```

**DESCRIPTION**

`nit_buf` is a STREAMS module that buffers incoming messages, thereby reducing the number of system calls and associated overhead required to read and process them. Although designed to be used in conjunction with the other components of NIT (see `nit(4P)`), `nit_buf` is a general-purpose module and can be used anywhere STREAMS input buffering is required.

**Read-side Behavior**

`nit_buf` collects incoming `M_DATA` and `M_PROTO` messages into *chunks*, passing each chunk upward when either the chunk becomes full or the current read timeout expires. When a message arrives, it is processed in two steps. First, the message is prepared for inclusion in a chunk, and then it is added to the current chunk. The following paragraphs discuss each step in turn.

Upon receiving a message from below, `nit_buf` immediately converts all leading `M_PROTO` blocks in the message to `M_DATA` blocks, altering only the message type field and leaving the contents alone. It then prepends a header to the converted message. This header is defined as follows.

```
struct nit_bufhdr {
    u_int  nhb_msglen;
    u_int  nhb_totlen;
};
```

The first field of this header gives the length in bytes of the converted message. The second field gives the distance in bytes from the start of the message in the current chunk (described below) to the start of the next message in the chunk; the value reflects any padding necessary to insure correct data alignment for the host machine and includes the length of the header itself.

After preparing a message, `nit_buf` attempts to add it to the end of the current chunk, using the chunk size and timeout values to govern the addition. (The chunk size and timeout values are set and inspected using the `ioctl` calls described below.) If adding the new message would make the current chunk grow larger than the chunk size, `nit_buf` closes off the current chunk, passing it up to the next module in line, and starts a new chunk, seeding it with a zero-length message. If adding the message would still make the current chunk overflow, the module passes it upward in an over-size chunk of its own. Otherwise, the module concatenates the message to the end of the current chunk.

To ensure that messages do not languish forever in an accumulating chunk, `nit_buf` maintains a read timeout. Whenever this timeout expires, the module closes off the current chunk, regardless of its length, and passes it upward; if no incoming messages have arrived, the chunk passed upward will have zero length. Whenever the module passes a chunk upward, it restarts the timeout period. These two rules insure that `nit_buf` minimizes the number of chunks it produces during periods of intense message activity and that it periodically disposes of all messages during slack intervals.

`nit_buf` handles other message types as follows. Upon receiving an `M_FLUSH` message specifying that the read queue be flushed, the module does so, clearing the currently accumulating chunk as well, and passes the message on to the module or driver above. It passes all other messages through unaltered to its upper neighbor.

**Write-side Behavior**

`nit_buf` intercepts `M_IOCTL` messages for the *ioctls* described below. Upon receiving an `M_FLUSH` message specifying that the write queue be flushed, the module does so and passes the message on to the module or driver below. The module passes all other messages through unaltered to its lower neighbor.

**IOCTLS**

**nit\_buf** responds to the following *ioctl*s.

- NIOCSTIME** Set the read timeout value to the value referred to by the *struct timeval* pointer given as argument. Setting the timeout value to zero has the side-effect of forcing the chunk size to zero as well, so that the module will pass all incoming messages upward immediately upon arrival.
- NIOCGTIME** Return the read timeout in the *struct timeval* pointed to by the argument. If the timeout has been cleared with the **NIOCCTIME** *ioctl*, return with an ERANGE error.
- NIOCCTIME** Clear the read timeout, effectively setting its value to infinity.
- NIOCSCHUNK** Set the chunk size to the value referred to by the *u\_int* pointer given as argument.
- NIOCGCHUNK** Return the chunk size in the *u\_int* pointed to by the argument.

**WARNING**

The module name “*nbuf*” used in the system configuration file and as argument to the **I\_PUSH** *ioctl* is provisional and subject to change.

**SEE ALSO**

**nit(4P)**, **nit\_if(4M)**, **nit\_pf(4M)**

**NAME**

`nit_if` – STREAMS NIT device interface module

**CONFIG**

`pseudo-device snit`

**SYNOPSIS**

```
#include <sys/file.h>
open("/dev/nit", mode);
```

**DESCRIPTION**

`nit_if` is a STREAMS pseudo-device driver that provides STREAMS access to network interfaces. It is designed to be used in conjunction with the other components of NIT (see `nit(4P)`), but can be used by itself as a raw STREAMS network interface.

`nit_if` is an exclusive-open device that is intended to be opened indirectly through the clone device; `/dev/nit` is a suitable instance of the clone device. Before the stream resulting from opening an instance of `nit_if` may be used to read or write packets, it must first be bound to a specific network interface, using the `NIOCSBIND` ioctl described below.

**Read-side Behavior**

`nit_if` copies leading prefixes of selected packets from its associated network interface and passes them up the stream. If the `NI_PROMISC` flag is set, it passes along all packets; otherwise it passes along only packets addressed to the underlying interface.

The amount of data copied from a given packet depends on the current *snapshot length*, which is set with the `NIOCSSNAP` ioctl described below.

Before passing each packet prefix upward, `nit_if` optionally prepends one or more headers, as controlled by the state of the flag bits set with the `NIOCSFLAGS` ioctl. The driver collects headers into `M_PROTO` message blocks, with the headers guaranteed to be completely contained in a single message block, whereas the packet itself goes into one or more `M_DATA` message blocks.

**Write-side Behavior**

`nit_if` accepts packets from the module above it in the stream and relays them to the associated network interface for transmission. Packets must be formatted with the destination address in a leading `M_PROTO` message block, followed by the packet itself, complete with link-level header, in a sequence of `M_DATA` message blocks. The destination address must be expressed as a `'struct sockaddr'` whose `sa_family` field is `AF_UNSPEC` and whose `sa_data` field is a copy of the link-level header. (See `sys/socket.h` for the definition of this structure.) If the packet does not conform to this format, an `M_ERROR` message with `EINVAL` will be sent upstream.

`nit_if` processes `M_IOCTL` messages as described below. Upon receiving an `M_FLUSH` message specifying that the write queue be flushed, `nit_if` does so and transfers the message to the read side of the stream. It discards all other messages.

**IOCTLS**

`nit_if` responds to the following *ioctls*, as defined in `net/nit_if.h`. It generates an `M_IOCNAK` message for all others, returning this message to the invoker along the read side of the stream.

**SIOCGIFADDR****SIOCADDMULTI****SIOCDELMULTI**

`nit_if` passes these *ioctls* on to the underlying interface's driver and returns its response in a `'struct ifreq'` instance, as defined in `net/if.h`. (See the description of this ioctl in `if(4N)` for more details.)

**NIOCBIND**

This ioctl attaches the stream represented by its first argument to the network interface designated by its third argument, which should be a pointer to an `ifreq` structure whose `ifr_name` field names the desired interface. See `net/if.h` for the definition of this structure.



- NIOCSSNAP** Set the current snapshot length to the value given in the *u\_long* pointed to by the *ioctl*'s final argument. *nit\_if* interprets a snapshot length value of zero as meaning infinity, so that it will copy all selected packets in their entirety. It constrains positive snapshot lengths to be at least the length of an Ethernet header, so that it will pass at least the link-level header of all selected packets to its upstream neighbor.
- NIOCGSNAP** Returns the current snapshot length for this device instance in the *u\_long* pointed to by the *ioctl*'s final argument.
- NIOCSFLAGS** *nit\_if* recognizes the following flag bits, which must be given in the *u\_long* pointed to by the *ioctl*'s final argument. This set may be augmented in future releases. All but the **NI\_PROMISC** bit control the addition of headers that precede the packet body. These headers appear in the order given below, with the last-mentioned enabled header adjacent to the packet body.
- NI\_PROMISC** Requests that the underlying interface be set into promiscuous mode and that all packets that the interface receives be passed up through the stream. *nit\_if* only honors this bit for the super-user.
- NI\_TIMESTAMP** Prepend to each selected packet a header containing the packet arrival time expressed as a 'struct timeval'.
- NI\_DROPS** Prepend to each selected packet a header containing the cumulative number of packets that this instance of *nit\_if* has dropped because of flow control requirements or resource exhaustion. The header value is expressed as a *u\_long*. Note: it accounts only for events occurring within *nit\_if*, and does not count packets dropped at the network interface level or by upstream modules.
- NI\_LEN** Prepend to each selected packet a header containing the packet's original length (including link-level header), as it was before being trimmed to the snapshot length. The header value is expressed as a *u\_long*.
- NIOCGFLAGS** Returns the current state of the flag bits for this device instance in the *u\_long* pointed to by the *ioctl*'s final argument.

**FILES**

- /dev/nit** clone device instance referring to *nit\_if* device
- net/nit\_if.h** header file containing definitions for the *ioctls* and packet headers described above.

**SEE ALSO**

**clone(4), nit(4P), nit\_buf(4M), nit\_pf(4M)**

**NAME**

`nit_pf` – STREAMS NIT packet filtering module

**CONFIG**

`pseudo-device pf`

**SYNOPSIS**

```
#include <sys/ioctl.h>
#include <net/nit_pf.h>
ioctl(fd, I_PUSH, "pf");
```

**DESCRIPTION**

`nit_pf` is a STREAMS module that subjects messages arriving on its read queue to a packet filter and passes only those messages that the filter accepts on to its upstream neighbor. Such filtering can be very useful for user-level protocol implementations and for networking monitoring programs that wish to view only specific types of events.

**Read-side Behavior**

`nit_pf` applies the current packet filter to all `M_DATA` and `M_PROTO` messages arriving on its read queue. The module prepares these messages for examination by first skipping over all leading `M_PROTO` message blocks to arrive at the beginning of the message's data portion. If there is no data portion, `nit_pf` accepts the message and passes it along to its upstream neighbor. Otherwise, the module ensures that the part of the message's data that the packet filter might examine lies in contiguous memory, calling the `pullupmsg` utility routine if necessary to force contiguity. (Note: this action destroys any sharing relationships that the subject message might have had with other messages.) Finally, it applies the packet filter to the message's data, passing the entire message upstream to the next module if the filter accepts, and discarding the message otherwise. See **PACKET FILTERS** below for details on how the filter works.

If there is no packet filter yet in effect, the module acts as if the filter exists but does nothing, implying that all incoming messages are accepted. **IOCTLS** below describes how to associate a packet filter with an instance of `nit_pf`.

`nit_pf` handles other message types as follows. Upon receiving an `M_FLUSH` message specifying that the read queue be flushed, the module does so, and passes the message on to its upstream neighbor. It passes all other messages through unaltered to its upper neighbor.

**Write-side Behavior**

`nit_pf` intercepts `M_IOCTL` messages for the `ioctl` described below. Upon receiving an `M_FLUSH` message specifying that the write queue be flushed, the module does so and passes the message on to the module or driver below. The module passes all other messages through unaltered to its lower neighbor.

**IOCTLS**

`nit_pf` responds to the following `ioctl`.

**NIOCSETF** This `ioctl` directs the module to replace its current packet filter, if any, with the filter specified by the 'struct `packetfilt`' pointer named by its final argument. This structure is defined in `<net/packetfilt.h>` as

```
struct packetfilt {
    u_char  Pf_Priority; /* priority of filter */
    u_char  Pf_FilterLen; /* # of cmds in list */
    u_short Pf_Filter[ENMAXFILTERS];
                                /* filter command list */
};
```

The *Pf\_Priority* field is included only for compatibility with other packet filter implementations and is otherwise ignored. The packet filter itself is specified in the *Pf\_Filter* array as a sequence of two-byte commands, with the *Pf\_FilterLen* field giving the number of commands in the sequence. This implementation restricts the maximum number of commands in a filter (ENMAXFILTERS) to 40. The next section describes the available commands and their semantics.

## PACKET FILTERS

A packet filter consists of the filter command list length (in units of *u\_shorts*), and the filter command list itself. (The priority field mentioned above is ignored in this implementation.) Each filter command list specifies a sequence of actions that operate on an internal stack of *u\_shorts* ("shortwords"). Each shortword of the command list specifies one of the actions ENF\_PUSHLIT, ENF\_PUSHZERO, or ENF\_PUSHPWORD+*n*, which respectively push the next shortword of the command list, zero, or shortword *n* of the subject message on the stack, and a binary operator from the set { ENF\_EQ, ENF\_NEQ, ENF\_LT, ENF\_LE, ENF\_GT, ENF\_GE, ENF\_AND, ENF\_OR, ENF\_XOR } which then operates on the top two elements of the stack and replaces them with its result. When both an action and operator are specified in the same shortword, the action is performed followed by the operation.

The binary operator can also be from the set { ENF\_COR, ENF\_CAND, ENF\_CNOR, ENF\_CNAND }. These are "short-circuit" operators, in that they terminate the execution of the filter immediately if the condition they are checking for is found, and continue otherwise. All pop two elements from the stack and compare them for equality; ENF\_CAND returns false if the result is false; ENF\_COR returns true if the result is true; ENF\_CNAND returns true if the result is false; ENF\_CNOR returns false if the result is true. Unlike the other binary operators, these four do not leave a result on the stack, even if they continue.

The short-circuit operators should be used when possible, to reduce the amount of time spent evaluating filters. When they are used, you should also arrange the order of the tests so that the filter will succeed or fail as soon as possible; for example, checking the IP destination field of a UDP packet is more likely to indicate failure than the packet type field.

The special action ENF\_NOPUSH and the special operator ENF\_NOP can be used to only perform the binary operation or to only push a value on the stack. Since both are (conveniently) defined to be zero, indicating only an action actually specifies the action followed by ENF\_NOP, and indicating only an operation actually specifies ENF\_NOPUSH followed by the operation.

After executing the filter command list, a non-zero value (true) left on top of the stack (or an empty stack) causes the incoming packet to be accepted and a zero value (false) causes the packet to be rejected. (If the filter exits as the result of a short-circuit operator, the top-of-stack value is ignored.) Specifying an undefined operation or action in the command list or performing an illegal operation or action (such as pushing a shortword offset past the end of the packet or executing a binary operator with fewer than two shortwords on the stack) causes a filter to reject the packet.

## EXAMPLES

The reverse ARP daemon program (*rarpd*(8C)) uses code similar to the following fragment to construct a filter that rejects all but RARP packets. That is, it accepts only packets whose Ethernet type field has the value ETHERTYPE\_REVARP.

```

struct ether_header eh;          /* used only for offset values */
struct packetfilt pf;
register u_short *fwp = pf.Pf_Filter;
u_short offset;

/*
 * Set up filter. Offset is the displacement of the Ethernet
 * type field from the beginning of the packet in units of
 * u_shorts.
*/

```

```

offset = ((u_int) &eh.ether_type - (u_int) &eh.ether_dhost) / sizeof (u_short);
*fwp++ = ENF_PUSHPWORD + offset;
*fwp++ = ENF_PUSHLIT;
*fwp++ = htons(ETHERTYPE_REVARP);
*fwp++ = ENF_EQ;
pf.Pf_FilterLen = fwp - &pf.Pf_Filter[0];

```

This filter can be abbreviated by taking advantage of the ability to combine actions and operations:

```

...
*fwp++ = ENF_PUSHPWORD + offset;
*fwp++ = ENF_PUSHLIT | ENF_EQ;
*fwp++ = htons(ETHERTYPE_REVARP);
...

```

#### WARNINGS

The module name 'pf' used in the system configuration file and as argument to the `I_PUSH ioctl` is provisional and subject to change.

The `Pf_Priority` field of the `packetfilt` structure is likely to be removed.

#### SEE ALSO

`inet(4F)`, `nit(4P)`, `nit_buf(4M)`, `nit_if(4M)`

**NAME**

null – data sink

**CONFIG**

None; included with standard system.

**SYNOPSIS**

```
#include <fcntl.h>
```

```
open("/dev/null", mode);
```

**DESCRIPTION**

Data written on the **null** special file is discarded.

Reads from the **null** special file always return an end-of-file indication.

**FILES**

/dev/null

**NAME**

openprom – PROM monitor configuration interface

**CONFIG**

pseudo-device openeepr

**SYNOPSIS**

```
#include <fcntl.h>
#include <sys/types.h>
#include <sundev/openpromio.h>
open("/dev/openprom", mode);
```

**AVAILABILITY**

SPARCstation 1 systems only.

**DESCRIPTION**

As with other Sun systems, configuration options are stored in an EEPROM or NVRAM on a SPARCstation 1 system. However, unlike other Sun systems, the encoding of these options is private to the PROM monitor. The `openprom` device provides an interface to the PROM monitor allowing a user program to query and set these configuration options through the use of `ioctl(2)` requests. These requests are defined in `<sundev/openpromio.h>`:

```
struct openpromio {
    u_int   oprom_size;           /* real size of following array */
    char    oprom_array[1];      /* For property names and values */
                                        /* NB: Adjacent, Null terminated */
};
#define OPROMMAXPARAM    1024    /* max size of array */

#define OPROMGETOPT      _IO(O,1)
#define OPROMSETOPT     _IO(O,2)
#define OPROMNXTOPT     _IO(O,3)
```

For all `ioctl()` requests, the third parameter is a pointer to a `'struct openpromio'`. All property names and values are null-terminated strings; the value of a numeric option is its ASCII representation.

**IOCTLS**

The `OPROMGETOPT` `ioctl` takes the null-terminated name of a property in the `oprom_array` and returns its null-terminated value (overlying its name). `oprom_size` should be set to the size of `oprom_array`; on return it will contain the size of the returned value. If the named property does not exist, or if there is not enough space to hold its value, then `oprom_size` will be set to zero. See **BUGS** below.

The `OPROMSETOPT` `ioctl` takes two adjacent strings in `oprom_array`; the null-terminated property name followed by the null-terminated value.

The `OPROMNXTOPT` `ioctl` is used to retrieve properties sequentially. The null-terminated name of a property is placed into `oprom_array` and on return it is replaced with the null-terminated name of the next property in the sequence, with `oprom_size` set to its length. A null string on input means return the name of the first property; an `oprom_size` of zero on output means there are no more properties.

**ERRORS**

<code>EINVAL</code>	The size value was invalid, or (for <code>OPROMSETOPT</code> ) the property does not exist.
<code>ENOMEM</code>	The kernel could not allocate space to copy the user's structure

**FILES**

<code>/dev/openprom</code>	PROM monitor configuration interface
----------------------------	--------------------------------------

**SEE ALSO**

`mem(4S)`, `eeprom(8S)`, `monitor(8S)`

**BUGS**

There should be separate return values for non-existent properties as opposed to not enough space for the value.

An attempt to set a property to an illegal value results in the PROM setting it to some legal value, with no error being returned. An **OPROMGETOPT** should be performed after an **OPROMSETOPT** to verify that the set worked.

The driver should be more consistent in its treatment of errors and edge conditions.

**NAME**

pp – Centronics-compatible parallel printer port

**CONFIG — Sun386i SYSTEMS**

device pp0 at obio ? csr 0x378 irq 15 priority 2

**CONFIG — SUN-3x SYSTEMS**

device pp0 at obio ? csr 0x6f000000 priority 1

This synopsis line should be used to generate a kernel for Sun-3/80 systems only.

**AVAILABILITY**

Sun386i and Sun-3/80 systems only.

**DESCRIPTION**

This device driver provides an interface to the Sun386i and Sun-3/80 systems' on-board Centronics-compatible parallel printer port. It supports most standard PC printers with Centronics interfaces.

**FILES**

/dev/pp0

**DIAGNOSTICS**

pp\*: printer not online

pp\*: printer out of paper



**NAME**

pty – pseudo-terminal driver

**CONFIG**

**pseudo-device** *ptyn*

**SYNOPSIS**

```
#include <fcntl.h>
#include <sys/termios.h>
open("/dev/ttypn", mode);
open("/dev/ptypn", mode);
```

**DESCRIPTION**

The **pty** driver provides support for a pair of devices collectively known as a *pseudo-terminal*. The two devices comprising a pseudo-terminal are known as a *controller* and a *slave*. The slave device distinguishes between the **B0** baud rate and other baud rates specified in the **c\_cflag** word of the **termios** structure, and the **CLOCAL** flag in that word. It does not support any of the other **termio(4)** device control functions specified by flags in the **c\_cflag** word of the **termios** structure and by the **IGNBRK**, **IGNPAR**, **PARMRK**, or **INPCK** flags in the **c\_iflag** word of the **termios** structure, as these functions apply only to asynchronous serial ports. All other **termio(4)** functions must be performed by **STREAMS** modules pushed atop the driver; when a slave device is opened, the **ldterm(4M)** and **ttcompat(4M)** **STREAMS** modules are automatically pushed on top of the stream, providing the standard **termio(4)** interface.

Instead of having a hardware interface and associated hardware that supports the terminal functions, the functions are implemented by another process manipulating the controller device of the pseudo-terminal.

The controller and the slave devices of the pseudo-terminal are tightly connected. Any data written on the controller device is given to the slave device as input, as though it had been received from a hardware interface. Any data written on the slave terminal can be read from the controller device (rather than being transmitted from a UART).

In configuring, if no optional “count” is given in the specification, 16 pseudo-terminal pairs are configured.

**IOCTLS**

The standard set of **termio** **ioctl**s are supported by the slave device. None of the bits in the **c\_cflag** word have any effect on the pseudo-terminal, except that if the baud rate is set to **B0**, it will appear to the process on the controller device as if the last process on the slave device had closed the line; thus, setting the baud rate to **B0** has the effect of “hanging up” the pseudo-terminal, just as it has the effect of “hanging up” a real terminal.

There is no notion of “parity” on a pseudo-terminal, so none of the flags in the **c\_iflag** word that control the processing of parity errors have any effect. Similarly, there is no notion of a “break”, so none of the flags that control the processing of breaks, and none of the **ioctl**s that generate breaks, have any effect.

Input flow control is automatically performed; a process that attempts to write to the controller device will be blocked if too much unconsumed data is buffered on the slave device. The input flow control provided by the **IXOFF** flag in the **c\_iflag** word is not supported.

The delays specified in the **c\_oflag** word are not supported.

As there are no modems involved in a pseudo-terminal, the **ioctl**s that return or alter the state of modem control lines are silently ignored.

On Sun systems, an additional **ioctl** is provided:

**TIOCCONS**

The argument is ignored. All output that would normally be sent to the console (either from programs writing to **/dev/console** or from kernel printouts) is redirected so that it is written to the pseudo-terminal instead.

A few special **ioctl**s are provided on the controller devices of pseudo-terminals to provide the functionality needed by applications programs to emulate real hardware interfaces:

#### **TIOCSTOP**

The argument is ignored. Output to the pseudo-terminal is suspended, as if a **STOP** character had been typed.

#### **TIOCSTART**

The argument is ignored. Output to the pseudo-terminal is restarted, as if a **START** character had been typed.

#### **TIOCPKT**

The argument is a pointer to an **int**. If the value of the **int** is non-zero, *packet* mode is enabled; if the value of the **int** is zero, packet mode is disabled. When a pseudo-terminal is in packet mode, each subsequent **read(2V)** from the controller device will return data written on the slave device preceded by a zero byte (symbolically defined as **TIOCPKT\_DATA**), or a single byte reflecting control status information. In the latter case, the byte is an inclusive-or of zero or more of the bits:

##### **TIOCPKT\_FLUSHREAD**

whenever the read queue for the terminal is flushed.

##### **TIOCPKT\_FLUSHWRITE**

whenever the write queue for the terminal is flushed.

##### **TIOCPKT\_STOP**

whenever output to the terminal is stopped using **^S**.

##### **TIOCPKT\_START**

whenever output to the terminal is restarted.

##### **TIOCPKT\_DOSTOP**

whenever **XON/XOFF** flow control is enabled after being disabled; it is considered "enabled" when the **IXON** flag in the **c\_iflag** word is set, the **VSTOP** member of the **c\_cc** array is **^S** and the **VSTART** member of the **c\_cc** array is **^Q**.

##### **TIOCPKT\_NOSTOP**

whenever **XON/XOFF** flow control is disabled after being enabled.

This mode is used by **rlogin(1C)** and **rlogind(8C)** to implement a remote-echoed, locally **^S/^Q** flow-controlled remote login with proper back-flushing of output when interrupts occur; it can be used by other similar programs.

#### **TIOCREMOTE**

The argument is a pointer to an **int**. If the value of the **int** is non-zero, *remote* mode is enabled; if the value of the **int** is zero, remote mode is disabled. This mode can be enabled or disabled independently of packet mode. When a pseudo-terminal is in remote mode, input to the slave device of the pseudo-terminal is flow controlled and not input edited (regardless of the mode the slave side of the pseudo-terminal). Each write to the controller device produces a record boundary for the process reading the slave device. In normal usage, a write of data is like the data typed as a line on the terminal; a write of 0 bytes is like typing an EOF character. Note: this means that a process writing to a pseudo-terminal controller in *remote* mode must keep track of line boundaries, and write only one line at a time to the controller. If, for example, it were to buffer up several **NEWLINE** characters and write them to the controller with one **write()**, it would appear to a process reading from the slave as if a single line containing several **NEWLINE** characters had been typed (as if, for example, a user had typed the **LNEXT** character before typing all but the last of those **NEWLINE** characters). Remote mode can be used when doing remote line editing in a window manager, or whenever flow controlled input is required.

The **ioctl**s **TIOCGWINSZ**, **TIOCSWINSZ**, and, on Sun systems, **TIOCCONS**, can be performed on the controller device of a pseudo-terminal; they have the same effect as when performed on the slave device.

**FILES**

**/dev/pty[p-s][0-9a-f]** pseudo-terminal controller devices  
**/dev/tty[p-s][0-9a-f]** pseudo-terminal slave devices  
**/dev/console**

**SEE ALSO**

**rlogin(1C), termio(4), ldterm(4M), ttcompat(4M), rlogind(8C)**

**BUGS**

It is apparently not possible to send an EOT by writing zero bytes in TIOCREMOTE mode.

**NAME**

rfs, RFS – remote file sharing

**CONFIGURATION**

**options RFS**  
**options VFSSTATS**

**AVAILABILITY**

Available only with the *RFS* software installation option. Refer to *Installing SunOS 4.1* for information on how to install optional software.

**DESCRIPTION**

The Remote File Sharing service, or RFS, allows transparent resource sharing among hosts on a network. A *resource* can be a directory, the files contained in that directory, subdirectories, devices, and even named pipes. Resources are advertised as a local directory using the name services. Hosts can then mount these resources, and use them as they would a local file system. The host advertising the resource is a file server, the hosts mounting the resource are clients.

All file servers and clients on a network belong to an RFS *domain*, and are administered by the same RFS name server. A domain consists of the following:

- A primary name server
- Possibly one or more secondary name servers
- File servers
- Clients

The name server maintains a list of advertised resources, and passwords in use. The name server also provides *name-to-resource* mapping. This allows a client to mount an advertised resource by the resource name, without needing to know the name of the file server or the pathname of the directory.

**FILES**

**/usr/nserve/rfmaster** hosts providing domain name service

**SEE ALSO**

**clone(4), nit\_buf(4M), nit\_pm(4M), tcptli(4P), timod(4), tirdwr(4), rfadmin(8), rfstart(8), rfdaemon(8), rmntstat(8)**

*System and Network Administration*

**NAME**

root – pseudo-driver for Sun386i root disk

**CONFIG**

**pseudo-device rootdev**

**AVAILABILITY**

Available only on Sun 386i systems running a SunOS 4.0.x release or earlier. Not a SunOS 4.1 release feature.

**DESCRIPTION**

The **root** pseudo-driver provides indirect, device-independent access to the root disk on a diskful Sun workstation. The root disk is the disk where the mounted root partition resides - typically the disk from which the system was booted.

The intent of the **root** device is to allow uniform access to the partitions on the root disk, regardless of the disk's controller type or unit number. For example, the following version of */etc/fstab* will work for any disk (assuming the disk has the standard partitions and filesystems):

```
/dev/roota / 4.2 rw 1 1
/dev/rootg /usr 4.2 ro 1 2
/dev/rootb /export 4.2 rw 1 3
```

When the root device is opened, the open and all subsequent operations on that device (**read(2V)**, **write(2V)**, **ioctl(2)**, **close(2V)**) are redirected to the real disk. Therefore, all device-dependent operations on a particular disk are still accessible via the root device (see **dkio(4S)**).

**FILES**

```
/dev/root[a-h]    block partitions
/dev/rroot[a-h]  raw partitions
```

**SEE ALSO**

**fstab(5)**, **sd(4S)**, **open(2V)**, **dkio(4S)**

**NAME**

routing – system supporting for local network packet routing

**DESCRIPTION**

The network facilities provided general packet routing, leaving routing table maintenance to applications processes.

A simple set of data structures comprise a “routing table” used in selecting the appropriate network interface when transmitting packets. This table contains a single entry for each route to a specific network or host. A user process, the routing daemon, maintains this data base with the aid of two socket specific `ioctl(2)` commands, `SIOCADDRT` and `SIOCDELRT`. The commands allow the addition and deletion of a single routing table entry, respectively. Routing table manipulations may only be carried out by super-user.

A routing table entry has the following form, as defined in `<net/route.h>`:

```
struct rtenry {
    u_long  rt_hash;
    struct  sockaddr rt_dst;
    struct  sockaddr rt_gateway;
    short   rt_flags;
    short   rt_refcnt;
    u_long  rt_use;
    struct  ifnet *rt_ifp;
};
```

with `rt_flags` defined from:

```
#define RTF_UP      0x1      /* route usable */
#define RTF_GATEWAY 0x2      /* destination is a gateway */
#define RTF_HOST    0x4      /* host entry (net otherwise) */
```

Routing table entries come in three flavors: for a specific host, for all hosts on a specific network, for any destination not matched by entries of the first two types (a wildcard route). When the system is booted, each network interface autoconfigured installs a routing table entry when it wishes to have packets sent through it. Normally the interface specifies the route through it is a “direct” connection to the destination host or network. If the route is direct, the transport layer of a protocol family usually requests the packet be sent to the same host specified in the packet. Otherwise, the interface may be requested to address the packet to an entity different from the eventual recipient (that is, the packet is forwarded).

Routing table entries installed by a user process may not specify the hash, reference count, use, or interface fields; these are filled in by the routing routines. If a route is in use when it is deleted (`rt_refcnt` is non-zero), the resources associated with it will not be reclaimed until all references to it are removed.

The routing code returns `EEXIST` if requested to duplicate an existing entry, `ESRCH` if requested to delete a non-existent entry, or `ENOBUFS` if insufficient resources were available to install a new route.

User processes read the routing tables through the `/dev/kmem` device.

The `rt_use` field contains the number of packets sent along the route. This value is used to select among multiple routes to the same destination. When multiple routes to the same destination exist, the least used route is selected.

A wildcard routing entry is specified with a zero destination address value. Wildcard routes are used only when the system fails to find a route to the destination host and network. The combination of wildcard routes and routing redirects can provide an economical mechanism for routing traffic.

**FILES**

`/dev/kmem`

**SEE ALSO**

`ioctl(2)`, `route(8C)`, `routed(8C)`

## NAME

sd – driver for SCSI disk devices

## CONFIG — SUN-3, SUN-3x, and SUN-4 SYSTEMS

controller si0 at vme24d16 ? csr 0x200000 priority 2 vector siintr 0x40

controller si0 at obio ? csr 0x140000 priority 2

disk sd0 at si0 drive 0 flags 0

disk sd1 at si0 drive 1 flags 0

disk sd2 at si0 drive 8 flags 0

disk sd3 at si0 drive 9 flags 0

disk sd4 at si0 drive 16 flags 0

disk sd6 at si0 drive 24 flags 0

controller sc0 at vme24d16 ? csr 0x200000 priority 2 vector scintr 0x40

disk sd0 at sc0 drive 0 flags 0

disk sd1 at sc0 drive 1 flags 0

disk sd2 at sc0 drive 8 flags 0

disk sd3 at sc0 drive 9 flags 0

disk sd4 at sc0 drive 16 flags 0

disk sd6 at sc0 drive 24 flags 0

The first two **controller** lines above specify the first and second SCSI host adapters for Sun-3, Sun-3x, and Sun-4 VME systems. The third **controller** line specifies the first and only SCSI host adapter on Sun-3/50 and Sun-3/60 systems.

The four lines following the **controller** specification lines define the available **disk** devices, **sd0** – **sd6**.

The **flags** field is used to specify the SCSI device type to the host adapter. **flags** must be set to 0 to identify **disk** devices.

The **drive** value is calculated using the formula:

$$8 * target + lun$$

where *target* is the SCSI target, and *lun* is the SCSI logical unit number.

The next configuration block, following **si0** and **si1** above, describes the configuration for the older **sc0** host adapter. It uses the same configuration description as the **si0** host adapter.

## CONFIG — SPARCsystem 330 and SUN-3/80 SYSTEMS

controller sm0 at obio ? csr 0xfa000000 priority 2

disk sd0 at sm0 drive 0 flags 0

disk sd1 at sm0 drive 1 flags 0

disk sd2 at sm0 drive 8 flags 0

disk sd3 at sm0 drive 9 flags 0

disk sd4 at sm0 drive 16 flags 0

disk sd6 at sm0 drive 24 flags 0

The SPARCsystem 330 and Sun-3/80 use an on-board SCSI host adapter, **sm0**. It follows the same rules as described above for the Sun-3, Sun-3x, and Sun-4 section.

## CONFIG — SUN-4/110 SYSTEM

controller sw0 at obio 2 csr 0xa000000 priority 2

disk sd0 at sw0 drive 0 flags 0

disk sd1 at sw0 drive 1 flags 0

disk sd2 at sw0 drive 8 flags 0

disk sd3 at sw0 drive 9 flags 0

disk sd4 at sw0 drive 16 flags 0

disk sd6 at sw0 drive 24 flags 0

The Sun-4/110 uses an on-board SCSI host adapter, `sw0`. It follows the same rules as described above for the Sun-3, and Sun-4 section.

#### CONFIG — SUN-3/E SYSTEM

```
controller se0 at vme24d16 ? csr 0x300000 priority 2 vector se_intr 0x40
disk sd0 at se0 drive 0 flags 0
disk sd1 at se0 drive 1 flags 0
disk sd2 at se0 drive 8 flags 0
disk sd3 at se0 drive 9 flags 0
```

The Sun-3/E uses a VME-based SCSI host adapter, `se0`. It follows the same rules as described above for the Sun-3 and Sun-4 section.

#### CONFIG — Sun386i

```
controller wds0 at obmem ? csr 0xFB000000 dmachan 7 irq 16 priority 2
disk sd0 at wds0 drive 0 flags 0
disk sd1 at wds0 drive 8 flags 0
disk sd2 at wds0 drive 16 flags 0
```

The Sun386i configuration follows the same rules described above under the Sun-3 and Sun-4 configuration section.

#### CONFIG — SPARCstation 1 SYSTEMS

```
device-driver esp
scsibus0 at esp
disk sd0 at scsibus0 target 3 lun 0
disk sd1 at scsibus0 target 1 lun 0
disk sd2 at scsibus0 target 2 lun 0
disk sd3 at scsibus0 target 0 lun 0
```

The SPARCstation 1 configuration files specify a device driver (`esp`), and a SCSI bus attached to that device driver, and then disks on that SCSI bus at the SCSI Target and Logical Unit addresses are specified.

#### DESCRIPTION

Files with minor device numbers 0 through 7 refer to various portions of drive 0. The standard device names begin with “`sd`” followed by the drive number and then a letter a-h for partitions 0-7 respectively. The character `?` stands here for a drive number in the range 0-6.

The block-files access the disk using the system’s normal buffering mechanism and are read and written without regard to physical disk records. There is also a “raw” interface that provides for direct transmission between the disk and the user’s read or write buffer. A single read or write call usually results in one I/O operation; raw I/O is therefore considerably more efficient when many bytes are transmitted. The names of the raw files conventionally begin with an extra ‘`r.`’

I/O requests (such as `lseek (2V)`) to the SCSI disk must have an offset that is a multiple of 512 bytes (`DEV_BSIZE`), or the driver returns an `EINVAL` error. If the transfer length is not a multiple of 512 bytes, the transfer count is rounded up by the driver.

#### Disk Support

This driver handles the Adaptec ACB-4000 disk controller for ST-506 drives, the Emulex MD21 disk controller for ESDI drives, and embedded, CCS-compatible SCSI disk drives.

On Sun386i and SPARCstation 1 systems, this driver supports the CDC Wren III half-height, and Wren IV full-height SCSI disk drives.

The type of disk drive is determined using the SCSI inquiry command and reading the volume label stored on block 0 of the drive. The volume label describes the disk geometry and partitioning; it must be present or the disk cannot be mounted by the system.



The **sd?a** partition is normally used for the root file system on a disk, the **sd?b** partition as a paging area (e.g. swap), and the **sd?c** partition for pack-pack copying. **sd?c** normally maps the entire disk and may also be used as the mount point for secondary disks in the system. The rest of the disk is normally the **sd?g** partition. For the primary disk, the user file system is located here.

**FILES**

**/dev/sd[0-6][a-h]**      block files  
**/dev/rsd[0-6][a-h]**      raw files

**SEE ALSO**

**dkio(4S)**, **directory(3V)**, **lseek(2V)**, **read(2V)**, **write(2V)**

Product Specification for Wren IV SCSI Model 94171

Product Specification for Wren III SCSI Model 94161

Product Specification for Wren III SCSI Model 94211

Emulex MD21 Disk Controller Programmer Reference Manual

Adaptec ACB-4000 Disk Controller OEM Manual

**DIAGNOSTICS****sd?: sdtimer: I/O request timeout**

A tape I/O operation has taken too long to complete. A device or host adapter failure may have occurred.

**sd?: sdtimer: can't abort request**

The driver is unable to find the request in the disconnect queue to notify the device driver that it has failed.

**sd?: no space for inquiry data****sd?: no space for disk label**

The driver was unable to get enough space for temporary storage. The driver is unable to open the disk device.

**sd?: <%s>**

The driver has found a SCSI disk device and opened it for the first time. The disk label is displayed to notify the user.

**sd?: SCSI bus failure**

A host adapter error was detected. The system may need to be rebooted.

**sd?: single sector I/O failed**

The driver attempted to recover from a transfer by writing each sector, one at a time, and failed. The disk needs to be reformatted to map out the new defect causing this error.

**sd?: retry failed****sd?: rezero failed**

A disk operation failed. The driver first tries to recover by retrying the command, if that fails, the driver rezeros the heads to cylinder 0 and repeats the retries. A failure of either the retry or rezero operations results in these warning messages; the error recovery operation continues until the retry count is exhausted. At that time a hard error is posted.

**sd?: request sense failed**

The driver was attempting to determine the cause of an I/O failure and was unable to get more information. This implies that the disk device may have failed.

**sd?: warning, abs. block %d has failed %d times**

The driver is warning the user that the specified block has failed repeatedly.

**sd?: block %d needs mapping****sd?: reassigning defective abs. block %d**

The specified block has failed repeatedly and may soon become an unrecoverable failure. If the driver does not map out the specified block automatically, it is recommend that the user correct the problem.

**sd?: reassign block failed**

The driver attempted to map out a block having excessive soft errors and failed. The user needs to run format and repair the disk.

**sd?%c: cmd how blk %d (rel. blk %d)****sense key(0x%x): %s, error code(0x%x): %s**

An I/O operation (**cmd**), encountered an error condition at absolute block (**blk %d**), partition (**sd?%c**), or relative block (**rel. block %d**). The error recovery operation (**how**) indicates whether it *retry*'ed, *restored*, or *failed*. The **sense key** and **error code** of the error are displayed for diagnostic purposes. The absolute **blk** of the the error is used for mapping out the defective block. The **rel. blk** is the block (sector) in error, relative to the beginning of the partition involved. This is useful for using **icheck(8)** to repair a damaged file structure on the disk.

**SPARCstation 1 Diagnostics**

The diagnostics for SPARCstation 1 are much like as above. Below are some additional diagnostics you might see on a SPARCstation 1:

**sd?: SCSI transport failed: reason 'xxxx': {retrying|giving up}**

The host adapter has failed to transport a command to the target for the reason stated. The driver will either retry the command or, ultimately, give up.

**sd?: disk not responding to selection**

The target disk isn't responding. You may have accidently kicked a power cord loose.

**sd?: disk ok**

The target disk is now responding again.

**sd?: disk offline**

The driver has decided that the target disk is no longer there.

**BUGS**

These disk drivers assume that you don't have removable media drives, and also that in order to operate normally, a valid Sun disk label must be in sector zero.

A logical block size of 512 bytes is assumed (and enforced on SPARCstation 1).

**NAME**

sockio – ioctls that operate directly on sockets

**SYNOPSIS**

```
#include <sys/sockio.h>
```

**DESCRIPTION**

The IOCTL's listed in this manual page apply directly to sockets, independent of any underlying protocol. Note: the **setsockopt** system call (see **getsockopt(2)**) is the primary method for operating on sockets as such, rather than on the underlying protocol or network interface. **ioctls** for a specific network interface or protocol are documented in the manual page for that interface or protocol.

- SIOCSGRP**           The argument is a pointer to an **int**. Set the process-group ID that will subsequently receive **SIGIO** or **SIGURG** signals for the socket referred to by the descriptor passed to **ioctl** to the value of that **int**.
- SIOCGRP**            The argument is a pointer to an **int**. Set the value of that **int** to the process-group ID that is receiving **SIGIO** or **SIGURG** signals for the socket referred to by the descriptor passed to **ioctl**.
- SIOCCATMARK**       The argument is a pointer to an **int**. Set the value of that **int** to 1 if the read pointer for the socket referred to by the descriptor passed to **ioctl** points to a mark in the data stream for an out-of-band message, and to 0 if it does not point to a mark.

**SEE ALSO**

**ioctl(2)**, **getsockopt(2)**, **filio(4)**

**NAME**

sr – driver for CDROM SCSI controller

**CONFIG — SPARCstation 1 and SPARCserver**

disk sr0 at scsibus0 target 6 lun 0

**CONFIG — SUN-4/330 SYSTEMS**

disk sr0 at sm0 drive 060 flags 2

**CONFIG — SUN-4 SYSTEMS**

disk sr0 at sc0 drive 060 flags 2

disk sr0 at si0 drive 060 flags 2

**AVAILABILITY**

SPARCstation 1, SPARCserver 1, and Sun-4/330 systems only.

**DESCRIPTION**

CDROM is a removable read-only direct-access device connected to the system's SCSI bus. CDROM drives are designed to work with any disc that meets the Sony-Philips "red-book" or "yellow-book" documents. They can read CDROM data discs, digital audio discs (Audio CD's) or combined-mode discs (that is, some tracks are audio, some tracks are data). A CDROM disc is singled sided containing approximately 540 mega-bytes of data or 74 minutes of audio.

The CDROM drive controller is set up as SCSI target 6. There is only a single logically unit number 0. Therefore, the minor device number is always 0.

Since all the other SCSI target ids has been reserved by the system, the system only supports one CDROM drive. The device names are `/dev/sr0` for block device and `/dev/rsr0` for character device.

The device driver supports `open(2V)`, `read(2V)`, `close(2V)` function calls through its block device and character device interface. In addition, it supports `ioctl` function call through the character device interface. When the device is first opened, the CDROM drive's eject button will be disabled (which prevents the manual removal of the disc) until the last `close(2V)` is called.

**CDROM Drive Support**

This driver supports the SONY CDU-8012 CDROM drive controller and other CDROM drives which has the same SCSI command set as the SONY CDU-8012. The type of CDROM drive is determined using the SCSI inquiry command.

There is no volume label stored on the CDROM. The disc geometry and partitioning information is always the same. If the CDROM is in ISO 9660 or High Sierra Disk format, it can be mounted as a file system.

**FILES**

<code>/dev/sr0</code>	block files
<code>/dev/rsr0</code>	raw files

**SEE ALSO**

`cdromio(4S)`, `fstab(5)`, `mount(8)`

## NAME

st – driver for SCSI tape devices

## CONFIG — SUN-3, SUN-3x, SUN-4 SYSTEMS

controller si0 at vme24d16 ? csr 0x200000 priority 2 vector siintr 0x40

controller si1 at vme24d16 ? csr 0x204000 priority 2 vector siintr 0x41

controller si0 at obio ? csr 0x140000 priority 2

tape st0 at si0 drive 32 flags 1

tape st1 at si0 drive 40 flags 1

tape st2 at si1 drive 32 flags 1

tape st3 at si1 drive 40 flags 1

controller sc0 at vme24d16 ? csr 0x200000 priority 2 vector scintr 0x40

tape st0 at sc0 drive 32 flags 1

tape st1 at sc0 drive 40 flags 1

The first two controller lines above specify the first and second SCSI host adapters for Sun-3, Sun-3x, and Sun-4 VME systems. The third controller line specifies the first and only SCSI host adapter on Sun-3/50 and Sun-3/60 systems.

Following the controller specification lines are four lines which define the available tape devices, st0–st3. The first two tape devices, st0 and st1, are on the first controller, si0. The next two tape devices, st2 and st3, are on the second controller, si1.

The flags field is used to specify the SCSI device type to the host adapter. The flags field must be set to 1 to identify tape devices.

The drive value is calculated using the formula:

$$8 * target + lun$$

where target is the SCSI target, and lun is the SCSI logical unit number.

The next configuration block, following si0 and si1 above, describes the older sc0 host adapter configuration. It follows the same configuration description as the si0 host adapter.

## CONFIG — SPARCsystem 330, SUN-3/80 SYSTEMS

controller sm0 at obio ? csr 0xfa000000 priority 2

tape st0 at sm0 drive 32 flags 1

tape st1 at sm0 drive 40 flags 1

The SPARCsystem 330 and Sun-3/80 use an on-board SCSI host adapter, sm0, which follows the rules described above in the Sun-3, Sun-3x, and Sun-4 section.

## CONFIG — SUN-4/110 SYSTEM

controller sw0 at obio 2 csr 0xa000000 priority 2

tape st0 at sw0 drive 32 flags 1

tape st1 at sw0 drive 40 flags 1

The Sun-4/110 uses an on-board SCSI host adapter, sw0, which follows the rules described above in the Sun-3, Sun-3x, and Sun-4 section.

## CONFIG — SUN-3/E SYSTEM

controller se0 at vme24d16 ? csr 0x300000 priority 2 vector se\_intr 0x40

tape st0 at se0 drive 32 flags 1

tape st1 at se0 drive 40 flags 1

The Sun-3/E uses a VME-based SCSI host adapter, se0, which follows the rules described above for Sun-3, Sun-3x, and Sun-4 systems.

**CONFIG — Sun386i**

**controller wds0 at obmem ? csr 0xFB000000 dmachan 7 irq 16 priority 2**  
**tape st0 at wds0 drive 32 flags 1**

The Sun386i configuration follows the rules described above in the Sun-3, Sun-3x, and Sun-4 configuration section.

**CONFIG — SPARCstation 1 SYSTEM**

**device-driver esp**  
**scsibus0 at esp**  
**tape st0 at scsibus0 target 4 lun 0**  
**tape st1 at scsibus0 target 5 lun 1**

The SPARCstation 1 configuration files specify a device driver (**esp**), and a SCSI bus attached to that device driver, and then tapes on that SCSI bus at the SCSI Target and Logical Unit addresses are specified.

**DESCRIPTION**

The **st** device driver is an interface to various SCSI tape devices. Supported 1/4-inch cartridge devices include the Archive Viper QIC-150 streaming tape drive, the Emulex MT-02 tape controller, and the Sysgen SC4000 (except on SPARCstation 1) tape controller. **st** provides a standard interface to these various devices, see **mtio(4)** for details.

The driver can be opened with either rewind on close (**/dev/rst\***) or no rewind on close (**/dev/nrst\***) options. A maximum of four tape formats per device are supported (see **FILES** below). The tape format is specified using the device name. The four rewind on close formats for **st0**, for example, are **/dev/rst0**, **/dev/rst8**, **/dev/rst16**, and **/dev/rst24**.

**Read Operation**

Fixed-length I/O tape devices require the number of bytes read or written to be a multiple of the physical record size. For example, 1/4-inch cartridge tape devices only read or write multiples of 512 bytes.

Fixed-length tape devices read or write multiple records if the blocking factor is greater than 64512 bytes (minphys limit). These multiple writes are limited to 64512 bytes. For example, if a write request is issued for 65536 bytes using a 1/4-inch cartridge tape, two writes are issued; the first for 64512 bytes and the second for 1024 bytes.

Tape devices, which support variable-length I/O operations, such as 1/2-inch reel tape, may read or write a range of 1 to 65535 bytes. If the record size exceeds 65535 bytes, the driver reads or writes multiple records to satisfy the request. These multiple records are limited to 65534 bytes. As an example, if a write request for 65540 bytes is issued using 1/2-inch reel tape, two records are written; one for 65534 bytes followed by one for 6 bytes.

If the driver is opened for reading in a different format than the tape is written in, the driver overrides the user selected format. For example, if a 1/4-inch cartridge tape is written in QIC-24 format and opened for reading in QIC-11, the driver will detect a read failure on the first read and automatically switch to QIC-24 to recover the data.

**Note:** If the **/dev/\*st[0-3]** format is used, no indication is given that the driver has overridden the user selected format. Other formats issue a warning message to inform the user of an overridden format selection. Some devices automatically perform this function and do not require driver support (1/2-inch reel and QIC-150 tape drives for example).

If a file mark is encountered during reading, no error is reported but the number of bytes transferred is zero. The next read operation reads into the next file.

End of media is indicated by two successive zero transfer counts. No further reading should be performed past the end of recorded media.

If the read request size is 2048 bytes, the tape driver behaves as a disk device and honors seek positioning requests (see **lseek(2)**). If a file mark is crossed during a read operation, this function is disabled.

**Write Operation**

Writing is allowed at either the beginning of tape or after the last written file on the tape. Writing from the beginning of tape is performed in the user-specified format. The original tape format is used for appending onto previously written tapes. A warning message is issued if the driver has to override the user-specified format.

Care should be used when appending files onto 1/2-inch reel tape devices, since an extra file mark is appended after the last file to mark the end of recorded media. In other words, the last file on the tape ends with two file marks instead of one. This extra file mark must be overwritten to prevent the creation of a null file. To facilitate write append operations, a space to the end of recorded media `ioctl()` is provided to eliminate this problem by having the driver perform the positioning operation.

If the end of tape is encountered during writing, no error is reported but the number of bytes transferred is zero and no further writing is allowed. Trailer records may be written by first writing a file mark followed by the trailer records. It is important that these trailer records be kept as short as possible to prevent data loss.

**Close Operation**

If data was written, a file mark is automatically written by the driver upon close. If the rewinding device name is used, the tape will be rewound after the file mark is written. If the user wrote a file mark prior to closing, then no file mark is written upon close. If a file positioning `ioctl()`, like `rewind`, is issued after writing, a file mark is written before repositioning the tape.

Note: For 1/2-inch reel tape devices, two file marks are written to mark the end of recorded media before rewinding or performing a file positioning `ioctl()`. If the user wrote mark before closing a 1/2-inch reel tape device, the driver will always write a file mark before closing to insure that the end of recorded media is marked properly.

If no data was written and the driver was opened for WRITE-ONLY access, a file mark is written thus creating a null file.

**IOCTLS**

The following `ioctl`s are supported: `forwardspace record`, `forwardspace file`, `backspace record`, `backspace file`, `backspace file mark`, `rewind`, `write file mark`, `offline`, `erase`, `retension`, `space to EOM`, and `get status`.

The `backspace file` and `forwardspace file` tape operations are inverses. Thus, a `forwardspace "-1"` file is equivalent to a `backspace "1"` file. A `backspace "0"` file is the same as `forwardspace "0"` file; both position the tape device to the beginning of the current file.

`Backspace file mark` moves the tape backwards by file marks. The tape position will end on the beginning of tape side of the desired file mark. Devices which do not support this function, such as 1/4-inch cartridge tape, return an `ENXIO` error.

`Backspace record` and `forwardspace record` operations perform much like `space file` operations, except that they move by records instead of files. Variable-length I/O devices (1/2-inch reel, for example) space actual records; fixed-length I/O devices space physical records (blocks). 1/4-inch cartridge tape, for example, spaces 512 byte physical records. The status `ioctl` residue count contains the number of files or records not skipped. Record skipping does not go past a file mark; file skipping does not go past the end of recorded media.

`Spacing to the end of recorded media` positions the tape at a location just after the last file written on the tape. For 1/4-inch cartridge tape, this is after the last file mark on the tape. For 1/2-inch reel tape, this is just after the first file mark but before the second (and last) file mark on the tape. Additional files can then be appended onto the tape from that point.

The `offline` `ioctl` rewinds and, if appropriate, takes the device offline by unloading the tape. Tape must be inserted before the tape device can be used again.

The `erase` `ioctl` rewinds the tape, erases it completely, and returns to the beginning of tape.

The `retension ioctl` only applies to 1/4-inch cartridge tape devices. It is used to restore tape tension improving the tape's soft error rate after extensive start-stop operations or long-term storage. Devices which do not support this function, such as 1/2-inch reel tape, return an `ENXIO` error.

The `get status ioctl` call returns the drive id (`mt_type`), sense key error (`mt_erreg`), file number (`mt_fileno`), and record number (`mt_blkno`) of the last error. The residue count (`mt_resid`) is set to the number of bytes not transferred or files/records not spaced.

Note: The error status is reset by the `get status ioctl` call or the next read, write, or other `ioctl` operation. If no error has occurred (sense key is zero), the current file and record position are returned.

## ERRORS

- EACCES** The driver is opened for write access and the tape is write protected, or an attempt is made to write on a write protected tape. For writing with QIC-150 tape drives, this error is also reported if the wrong tape media is used for writing.
- EBUSY** The tape device is already in use.
- EIO** During opening, the tape device is not ready because either no tape is in the drive, or the drive is not on-line. Once open, this error is returned if the requested I/O transfer could not be completed.
- EINVAL** The number of bytes read or written is not a multiple of the physical record size (fixed-length tape devices only).
- ENXIO** During opening, the tape device does not exist. On `ioctl` functions, this indicates that the tape device does not support the `ioctl` function.

## FILES

For QIC-150 tape devices (Archive Viper):

```

/dev/rst[0-3]   QIC-150 Format
/dev/rst[8-11]  QIC-150 Format
/dev/rst[16-20] QIC-150 Format
/dev/rst[24-28] QIC-150 Format
/dev/nrst[0-3]  non-rewinding QIC-150 Format
/dev/nrst[8-11] non-rewinding QIC-150 Format
/dev/nrst[16-19] non-rewinding QIC-150 Format
/dev/nrst[24-27] non-rewinding QIC-150 Format

```

For QIC-24 tape devices (Emulex MT-02 and Sysgen SC4000):

```

/dev/rst[0-3]   QIC-11 Format
/dev/rst[8-11]  QIC-24 Format
/dev/rst[16-20] QIC-24 Format
/dev/rst[24-28] QIC-24 Format
/dev/nrst[0-3]  non-rewinding QIC-11 Format
/dev/nrst[8-11] non-rewinding QIC-24 Format
/dev/nrst[16-19] non-rewinding QIC-24 Format
/dev/nrst[24-27] non-rewinding QIC-24 Format

```

Note: The QIC-24 format is preferred over QIC-11 for Sun-3, Sun-3x, Sun-4, and Sun386i systems.

## SEE ALSO

`mt(1)`, `tar(1)`, `mtio(4)`, `dump(8)`, `restore(8)`

Archive Viper QIC-150 Tape Drive Product Specification  
 Emulex MT-02 Intelligent Tape Controller Product Specification  
 Sysgen SC4000 Intelligent Tape Controller Product Specification



## DIAGNOSTICS

**st?: sttimer: I/O request timeout**

A tape I/O operation has taken too long to complete. A device or host adapter failure may have occurred.

**st?: sttimer: can't abort request**

The driver is unable to find the request in the disconnect que to notify the device driver that it has failed. A SCSI bus reset is issued to recover from this error.

**st?: unknown SCSI device found**

The SCSI device is not a tape device; it is some other type of SCSI device.

**st?: warning, unknown tape drive found**

The driver does not recognize the tape device. Only the default tape density is used; block size is set to the value specified by the tape drive.

**st?: tape is write protected**

The tape is write protected.

**st?: wrong tape media for writing**

For QIC-150 tape drives, this indicates that the user is trying to write on a DC-300XL (or equivalent) tape. Only DC-6150 (or equivalent) tapes can be used for writing.

Note: DC-6150 was formerly known as DC-600XTD.

**st?: warning, rewinding tape**

The driver is rewinding tape in order to set the tape format.

**st?: warning, using alternate tape format**

The driver is overriding the user-selected tape format and using the previously used format.

**st?: warning, tape rewound**

For Sysgen tape controllers, the tape may be rewound as a result of getting sense data.

**st?: format change failed**

The tape drive rejected the mode select command to change the tape format.

**st?: file mark write failed**

The driver was unable to write a file mark.

**st?: warning, The tape may be wearing out or the head may need cleaning.****st?: read retries= %d, file= %d, block= %d****st?: write retries= %d, file= %d, block= %d**

The number of allowable soft errors has been exceeded for this tape. Either the tape heads need cleaning or the tape is wearing out. If the tape is wearing out, continued usage of it is not recommended.

**st?: illegal command**

The SCSI command just issued was illegal. This message can result from issuing an inappropriate command, such as trying to write over previously written files on the tape. On foreign tape devices, this can also be caused by selecting the wrong tape format.

**st?: error: sense key(0x%x): %s, error code(0x%x): %s**

An error has occurred. The sense key message and error code are displayed for diagnostic purposes.

**st?: stread: not modulo %d block size****st?: stwrite: not modulo %d block size**

The read or write request size must be a multiple of the %d physical block size.

**st?: file positioning error****st?: block positioning error**

The driver was unable to position the tape to the desired file or block (record). This is probably caused by a damaged tape.

**st?: SCSI transport failed: reason 'xxxx': {retrying|giving up}**

The host adapter has failed to transport a command to the target for the reason stated. The driver will either retry the command or, ultimately, give up (SPARCstation 1) only.

**BUGS**

Foreign tape devices which do not return a BUSY status during tape loading prevent user commands from being held until the device is ready. The user must delay issuing any tape operations until the tape device is ready. This is not a problem for Sun supplied tape devices.

Foreign tape devices which do not report a blank check error at the end of recorded media cause file positioning operations to fail. Some tape drives for example, mistakenly report media error instead of blank check error.

“Cooked” mode for read and write operations is not supported.

Systems using the older `sc0` host adapter or the Sysgen SC4000 tape controller, prevent disk I/O over the SCSI bus while the tape is in use (during a rewind for example). This problem is caused by the fact that they do not support disconnect/reconnect to free the SCSI bus. Newer tape devices, like the the Emulex MT-02, and host adapters, like `si0`, eliminate this problem.

Some older systems may not support the QIC-24 format, and may complain (or exhibit erratic behavior) when the user attempts to use this format.

SPARCstation 1 does not support the Sysgen SC4000 tape controller, nor does it support 1/2" variable record length operations, record space operations, or implied seeking.

**NAME**

streamio – STREAMS ioctl commands

**SYNOPSIS**

```
#include <stropts.h>
int ioctl (fd, command, arg)
int fd, command;
```

**DESCRIPTION**

STREAMS (see [intro\(2\)](#)) ioctl commands are a subset of [ioctl\(2\)](#) commands that perform a variety of control functions on STREAMS. The arguments *command* and *arg* are passed to the file designated by *fd* and are interpreted by the *streamhead*. Certain combinations of these arguments may be passed to a module or driver in the stream.

*fd* is an open file descriptor that refers to a stream. *command* determines the control function to be performed as described below. *arg* represents additional information that is needed by this command. The type of *arg* depends upon the command, but it is generally an integer or a pointer to a *command*-specific data structure.

Since these STREAMS commands are a subset of *ioctl*, they are subject to the errors described there. In addition to those errors, the call will fail with *errno* set to EINVAL, without processing a control function, if the stream referenced by *fd* is linked below a multiplexor, or if *command* is not a valid value for a *stream*.

Also, as described in *ioctl*, STREAMS modules and drivers can detect errors. In this case, the module or driver sends an error message to the *stream head* containing an error value. Subsequent system calls will fail with *errno* set to this value.

**IOCTLS**

The following *ioctl* commands, with error values indicated, are applicable to all STREAMS files:

**I\_PUSH** Pushes the module whose name is pointed to by *arg* onto the top of the current stream, just below the *streamhead*. It then calls the open routine of the newly-pushed module.

I\_PUSH will fail if one of the following occurs:

EINVAL	The module name is invalid.
EFAULT	<i>arg</i> points outside the allocated address space.
ENXIO	The open routine of the new module failed.
ENXIO	A hangup is received on the stream referred to by <i>fd</i> .

**I\_POP** Removes the module just below the *stream head* of the stream pointed to by *fd*. *arg* should be 0 in an I\_POP request.

I\_POP will fail if one of the following occurs:

EINVAL	No module is present on <i>stream</i> .
ENXIO	A hangup is received on the stream referred to by <i>fd</i> .

**I\_LOOK** Retrieves the name of the module just below the *stream head* of the stream pointed to by *fd*, and places it in a null-terminated character string pointed at by *arg*. The buffer pointed to by *arg* should be at least FMNAMESZ+1 bytes long. An '#include <sys/conf.h>' declaration is required.

I\_LOOK will fail if one of the following occurs:

EFAULT	<i>arg</i> points outside the allocated address space of the process.
EINVAL	No module is present on <i>stream</i> .

**I\_FLUSH**

This request flushes all input and/or output queues, depending on the value of *arg*. Legal *arg* values are:

**FLUSHR** Flush read queues.  
**FLUSHW** Flush write queues.  
**FLUSHRW** Flush read and write queues.

**I\_FLUSH** will fail if one of the following occurs:

**EAGAIN** No buffers could be allocated for the flush message.  
**EINVAL** The value of *arg* is invalid.  
**ENXIO** A hangup is received on the stream referred to by *fd*.

**I\_SETSIG**

Informs the *stream head* that the user wishes the kernel to issue the SIGPOLL signal (see sigvec(2)) when a particular event has occurred on the stream associated with *fd*. **I\_SETSIG** supports an asynchronous processing capability in STREAMS. The value of *arg* is a bitmask that specifies the events for which the user should be signaled. It is the bitwise-OR of any combination of the following constants:

**S\_INPUT** A non-priority message has arrived on a *stream head* read queue, and no other messages existed on that queue before this message was placed there. This is set even if the message is of zero length.  
**S\_HIPRI** A priority message is present on the *stream head* read queue. This is set even if the message is of zero length.  
**S\_OUTPUT** The write queue just below the *stream head* is no longer full. This notifies the user that there is room on the queue for sending (or writing) data downstream.  
**S\_MSG** A STREAMS signal message that contains the SIGPOLL signal has reached the front of the *stream head* read queue.

A user process may choose to be signaled only of priority messages by setting the *arg* bitmask to the value **S\_HIPRI**.

Processes that wish to receive SIGPOLL signals must explicitly register to receive them using **I\_SETSIG**. If several processes register to receive this signal for the same event on the same *stream*, each process will be signaled when the event occurs.

If the value of *arg* is zero, the calling process will be unregistered and will not receive further SIGPOLL signals.

**I\_SETSIG** will fail if one of the following occurs:

**EINVAL** The value of *arg* is invalid or *arg* is zero and the process is not registered to receive the SIGPOLL signal.  
**EAGAIN** A data structure could not be allocated to store the signal request.

**I\_GETSIG**

Returns the events for which the calling process is currently registered to be sent a SIGPOLL signal. The events are returned as a bitmask pointed to by *arg*, where the events are those specified in the description of **I\_SETSIG** above.

**I\_GETSIG** will fail if one of the following occurs:

- EINVAL**                   The process is not registered to receive the **SIGPOLL** signal.
- EFAULT**                   *arg* points outside the allocated address space of the process.

**I\_FIND**

This request compares the names of all modules currently present in the stream to the name pointed to by *arg*, and returns 1 if the named module is present in the stream. It returns 0 if the named module is not present.

**I\_FIND** will fail if one of the following occurs:

- EFAULT**                   *arg* points outside the allocated address space of the process.
- EINVAL**                   *arg* does not point to a valid module name.

**I\_PEEK**

This request allows a user to retrieve the information in the first message on the *stream head* read queue without taking the message off the queue. *arg* points to a *strpeek* structure which contains the following members:

```

    struct strbuf  ctlbuf;
    struct strbuf  databuf;
    long          flags;

```

The *maxlen* field in the *ctlbuf* and *databuf* *strbuf* structures (see `getmsg(2)`) must be set to the number of bytes of control information and/or data information, respectively, to retrieve. If the user sets *flags* to **RS\_HIPRI**, **I\_PEEK** will only look for a priority message on the *stream head* read queue.

**I\_PEEK** returns 1 if a message was retrieved, and returns 0 if no message was found on the *stream head* read queue, or if the **RS\_HIPRI** flag was set in *flags* and a priority message was not present on the *stream head* read queue. It does not wait for a message to arrive. On return, *ctlbuf* specifies information in the control buffer, *databuf* specifies information in the data buffer, and *flags* contains the value 0 or **RS\_HIPRI**.

**I\_PEEK** will fail if one of the following occurs:

- EFAULT**                   *arg* points, or the buffer area specified in *ctlbuf* or *databuf* is, outside the allocated address space of the process.

**I\_SRDOPT**

Sets the read mode using the value of the argument *arg*. Legal *arg* values are:

- RNORM**                   Byte-stream mode, the default.
- RMSGD**                   Message-discard mode.
- RMSGN**                   Message-nondiscard mode.

Read modes are described in `read(2V)`.

**I\_SRDOPT** will fail if one of the following occurs:

- EINVAL**                   *arg* is not one of the above legal values.

**I\_GRDOPT**

Returns the current read mode setting in an *int* pointed to by the argument *arg*. Read modes are described in `read(2V)`.

**I\_GRDOPT** will fail if one of the following occurs:

- EFAULT**                   *arg* points outside the allocated address space of the process.

**I\_NREAD**

Counts the number of data bytes in data blocks in the first message on the *stream head* read queue, and places this value in the location pointed to by *arg*. The return value for the command is the number of messages on the *stream head* read queue. For example, if zero is returned in *arg*, but the *ioctl* return value is greater than zero, this indicates that a zero-length message is next on the queue.

**I\_NREAD** will fail if one of the following occurs:

**EFAULT** *arg* points outside the allocated address space of the process.

**I\_FDINSERT**

creates a message from user specified buffer(s), adds information about another stream and sends the message downstream. The message contains a control part and an optional data part. The data and control parts to be sent are distinguished by placement in separate buffers, as described below.

*arg* points to a *strfdinsert* structure which contains the following members:

```

struct strbuf   ctlbuf;
struct strbuf   databuf;
long           flags;
int           fd;
int           offset;

```

The *len* field in the *ctlbuf strbuf* structure (see *putmsg(2)*) must be set to the size of a pointer plus the number of bytes of control information to be sent with the message. *fd* specifies the file descriptor of the other stream and *offset*, which must be word-aligned, specifies the number of bytes beyond the beginning of the control buffer where **I\_FDINSERT** will store a pointer to the *fd* stream's driver read queue structure. The *len* field in the *databuf strbuf* structure must be set to the number of bytes of data information to be sent with the message or zero if no data part is to be sent.

*flags* specifies the type of message to be created. A non-priority message is created if *flags* is set to 0, and a priority message is created if *flags* is set to **RS\_HIPRI**. For non-priority messages, **I\_FDINSERT** will block if the stream write queue is full due to internal flow control conditions. For priority messages, **I\_FDINSERT** does not block on this condition. For non-priority messages, **I\_FDINSERT** does not block when the write queue is full and **O\_NDELAY** is set. Instead, it fails and sets *errno* to **EAGAIN**.

**I\_FDINSERT** also blocks, unless prevented by lack of internal resources, waiting for the availability of message blocks in the *stream*, regardless of priority or whether **O\_NDELAY** has been specified. No partial message is sent.

**I\_FDINSERT** will fail if one of the following occurs:

**EAGAIN** A non-priority message was specified, the **O\_NDELAY** flag is set, and the stream write queue is full due to internal flow control conditions.

**EAGAIN** Buffers could not be allocated for the message that was to be created.

**EFAULT** *arg* points, or the buffer area specified in *ctlbuf* or *databuf* is, outside the allocated address space of the process.

EINVAL	<i>fd</i> in the <i>strfdinsert</i> structure is not a valid, open stream file descriptor; the size of a pointer plus <i>offset</i> is greater than the <i>len</i> field for the buffer specified through <i>ctlptr</i> ; <i>offset</i> does not specify a properly-aligned location in the data buffer; an undefined value is pointed to by <i>flags</i> .
ENXIO	A hangup is received on the stream referred to by <i>fd</i> .
ERANGE	The <i>len</i> field for the buffer specified through <i>databuf</i> does not fall within the range specified by the maximum and minimum packet sizes of the topmost stream module, or the <i>len</i> field for the buffer specified through <i>databuf</i> is larger than the maximum configured size of the data part of a message, or the <i>len</i> field for the buffer specified through <i>ctlbuf</i> is larger than the maximum configured size of the control part of a message.

**I\_STR**

Constructs an internal STREAMS ioctl message from the data pointed to by *arg*, and sends that message downstream.

This mechanism is provided to permit a process to specify timeouts and variable-sized amounts of data when sending an *ioctl* request to downstream modules and drivers. It allows information to be sent with the *ioctl*, and will return to the user any information sent upstream by the downstream recipient. **I\_STR** blocks until the system responds with either a positive or negative acknowledgement message, or until the request "times out" after some period of time. If the request times out, it fails with *errno* set to ETIME.

At most, one **I\_STR** can be active on a stream. Further **I\_STR** calls will block until the active **I\_STR** completes at the *stream head*. The default timeout interval for these requests is 15 seconds. The **O\_NDELAY** (see **open(2V)**) flag has no effect on this call.

To send requests downstream, *arg* must point to a *strioc* structure which contains the following members:

```

int    ic_cmd;        /* downstream command */
int    ic_timeout;    /* ACK/NAK timeout */
int    ic_len;        /* length of data arg */
char   *ic_dp;        /* ptr to data arg */

```

*ic\_cmd* is the internal *ioctl* command intended for a downstream module or driver and *ic\_timeout* is the number of seconds (-1 = infinite, 0 = use default, >0 = as specified) an **I\_STR** request will wait for acknowledgement before timing out. *ic\_len* is the number of bytes in the data argument and *ic\_dp* is a pointer to the data argument. The *ic\_len* field has two uses: on input, it contains the length of the data argument passed in, and on return from the command, it contains the number of bytes being returned to the user (the buffer pointed to by *ic\_dp* should be large enough to contain the maximum amount of data that any module or the driver in the stream can return).

The *stream head* will convert the information pointed to by the *strioc* structure to an internal *ioctl* command message and send it downstream.

**I\_STR** will fail if one of the following occurs:

EAGAIN	Buffers could not be allocated for the <i>ioctl</i> message.
--------	--

EFAULT	<i>arg</i> points, or the buffer area specified by <i>ic_dp</i> and <i>ic_len</i> (separately for data sent and data returned) is, outside the allocated address space of the process.
EINVAL	<i>ic_len</i> is less than 0 or <i>ic_len</i> is larger than the maximum configured size of the data part of a message or <i>ic_timeout</i> is less than -1.
ENXIO	A hangup is received on the stream referred to by <i>fd</i> .
ETIME	A downstream <i>ioctl</i> timed out before acknowledgement was received.

An *I\_STR* can also fail while waiting for an acknowledgement if a message indicating an error or a hangup is received at the *streamhead*. In addition, an error code can be returned in the positive or negative acknowledgement message, in the event the *ioctl* command sent downstream fails. For these cases, *I\_STR* will fail with *errno* set to the value in the message.

**I\_SENDFD**

Requests the stream associated with *fd* to send a message, containing a file pointer, to the *stream head* at the other end of a stream pipe. The file pointer corresponds to *arg*, which must be an integer file descriptor.

*I\_SENDFD* converts *arg* into the corresponding system file pointer. It allocates a message block and inserts the file pointer in the block. The user id and group id associated with the sending process are also inserted. This message is placed directly on the read queue (see *intro(2)*) of the *stream head* at the other end of the stream pipe to which it is connected.

*I\_SENDFD* will fail if one of the following occurs:

EAGAIN	The sending stream is unable to allocate a message block to contain the file pointer.
EAGAIN	The read queue of the receiving <i>stream head</i> is full and cannot accept the message sent by <i>I_SENDFD</i> .
EBADF	<i>arg</i> is not a valid, open file descriptor.
EINVAL	<i>fd</i> is not connected to a stream pipe.
ENXIO	A hangup is received on the stream referred to by <i>fd</i> .

**I\_RECVFD**

Retrieves the file descriptor associated with the message sent by an *I\_SENDFD* *ioctl* over a stream pipe. *arg* is a pointer to a data buffer large enough to hold an *strrecvfd* data structure containing the following members:

```
int fd;
unsigned short uid;
unsigned short gid;
char fill[8];
```

*fd* is an integer file descriptor. *uid* and *gid* are the user ID and group ID, respectively, of the sending stream.

If *O\_NDELAY* is not set (see *open(2V)*), *I\_RECVFD* will block until a message is present at the *streamhead*. If *O\_NDELAY* is set, *I\_RECVFD* will fail with *errno* set to *EAGAIN* if no message is present at the *streamhead*.

If the message at the *stream head* is a message sent by an *I\_SENDFD*, a new user file descriptor is allocated for the file pointer contained in the message. The new file descriptor is placed in the *fd* field of the *strrecvfd* structure. The structure is copied into the user data buffer pointed to by *arg*.



**I\_RECVFD** will fail if one of the following occurs:

EAGAIN	A message was not present at the <i>stream head</i> read queue, and the <b>O_NDELAY</b> flag is set.
EBADMSG	The message at the <i>stream head</i> read queue was not a message containing a passed file descriptor.
EFAULT	<i>arg</i> points outside the allocated address space of the process.
EMFILE	Too many descriptors are active.
ENXIO	A hangup is received on the stream referred to by <i>fd</i> .

The following four commands are used for connecting and disconnecting multiplexed STREAMS configurations.

### **I\_LINK**

Connects two streams, where *fd* is the file descriptor of the stream connected to the multiplexing driver, and *arg* is the file descriptor of the stream connected to another driver. The stream designated by *arg* gets connected below the multiplexing driver. **I\_LINK** causes the multiplexing driver to send an acknowledgement message to the *stream head* regarding the linking operation. This call returns a multiplexor ID number (an identifier used to disconnect the multiplexor, see **I\_UNLINK**) on success, and a -1 on failure.

**I\_LINK** will fail if one of the following occurs:

ENXIO	A hangup is received on the stream referred to by <i>fd</i> .
ETIME	The <b>ioctl</b> timed out before an acknowledgement was received.
EAGAIN	Storage could not be allocated to perform the <b>I_LINK</b> .
EBADF	<i>arg</i> is not a valid, open file descriptor.
EINVAL	The stream referred to by <i>fd</i> does not support multiplexing.
EINVAL	<i>arg</i> is not a stream, or is already linked under a multiplexor.
EINVAL	The specified link operation would cause a "cycle" in the resulting configuration; that is, if a given <i>stream head</i> is linked into a multiplexing configuration in more than one place.

An **I\_LINK** can also fail while waiting for the multiplexing driver to acknowledge the link request, if a message indicating an error or a hangup is received at the *stream head* of *fd*. In addition, an error code can be returned in the positive or negative acknowledgement message. For these cases, **I\_LINK** will fail with *errno* set to the value in the message.

### **I\_UNLINK**

Disconnects the two streams specified by *fd* and *arg*. *fd* is the file descriptor of the stream connected to the multiplexing driver. *arg* is the multiplexor ID number that was returned by the **ioctl** **I\_LINK** command when a stream was linked below the multiplexing driver. If *arg* is -1, then all streams which were linked to *fd* are disconnected. As in **I\_LINK**, this command requires the multiplexing driver to acknowledge the unlink.

**I\_UNLINK** will fail if one of the following occurs:

ENXIO	A hangup is received on the stream referred to by <i>fd</i> .
-------	---

- ETIME**                   The **ioctl** timed out before an acknowledgement was received.
- EAGAIN**                   Buffers could not be allocated for the acknowledgement message.
- EINVAL**                   The multiplexor ID number was invalid.

An **I\_UNLINK** can also fail while waiting for the multiplexing driver to acknowledge the link request, if a message indicating an error or a hangup is received at the *stream head* of *fd*. In addition, an error code can be returned in the positive or negative acknowledgement message. For these cases, **I\_UNLINK** will fail with *errno* set to the value in the message.

**SEE ALSO**

**close(2V)**, **fcntl(2V)**, **getmsg(2)**, **intro(2)**, **ioctl(2)**, **open(2V)**, **poll(2)**, **putmsg(2)**, **read(2V)**, **sigvec(2)**, **write(2V)**

*STREAMS Programmer's Guide*  
*STREAMS Primer*

**NAME**

taac – Sun applications accelerator

**CONFIG**

taac0 at vme32d32 ? csr 0x28000000

**CONFIG – SUN-3/SUN-4 SYSTEMS**

device taac0 at vme32d32 1 csr 0x28000000

device taac0 at vme32d32 2 csr 0xf8000000

device taac0 at vme32d32 3 csr 0x28000000

The first line should be used to generate a kernel for Sun-3/160, Sun-3/260, Sun-4/260, Sun-4/370 and Sun-4/460 systems. The second line should be used to generate a kernel for Sun-4/110 systems; and the last line should be used to generate a kernel for Sun-4/330 systems.

**CONFIG – SUN-4/150 SYSTEMS**

device taac0 at vme32d32 2 csr 0xf8000000

**AVAILABILITY**

TAAC-1 can only be used in Sun VME-bus packages with 4 or more full size (9U) slots.

**DESCRIPTION**

The taac interface supports the optional TAAC-1 Applications Accelerator. This add-on device is composed of a very-long-instruction-word computation engine, coupled with an 8MB memory array. This memory area can be used as a frame buffer or as storage for large data sets.

the Sun-4/150 VME address space is limited to 28 bits. The TAAC-1 must be reconfigured to work in this package. See *Configuration Procedures for the TAAC-1 Application Accelerator Board Set*.

Programs can be downloaded for execution on the TAAC-1 directly, they can be executed by the host processor, or the host processor and the TAAC-1 engine can be used in combination. See the *TAAC-1 User's Guide* for detailed information on accessing the TAAC-1 from the host. This manual also describes the C compiler, the programming tools, and the support libraries for the TAAC-1.

Programs on the host processor gain access to the TAAC-1 registers and memory by using `mmap(2)`.

**SEE ALSO**

`mmap(2)`

*TAAC-1 Application Accelerator: User Guide*

*Configuration Procedures for the TAAC-1 Application Accelerator Board Set*

**NAME**

tcp – Internet Transmission Control Protocol

**SYNOPSIS**

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

s = socket(AF_INET, SOCK_STREAM, 0);
```

**DESCRIPTION**

TCP is the virtual circuit protocol of the Internet protocol family. It provides reliable, flow-controlled, in order, two-way transmission of data. It is a byte-stream protocol used to support the **SOCK\_STREAM** abstraction. TCP is layered above the Internet Protocol (IP), the Internet protocol family's unreliable inter-network datagram delivery protocol.

TCP uses IP's host-level addressing and adds its own per-host collection of "port addresses". The endpoints of a TCP connection are identified by the combination of an IP address and a TCP port number. Although other protocols, such as the User Datagram Protocol (UDP), may use the same host and port address format, the port space of these protocols is distinct. See **inet(4F)** for details on the common aspects of addressing in the Internet protocol family.

Sockets utilizing TCP are either "active" or "passive". Active sockets initiate connections to passive sockets. Both types of sockets must have their local IP address and TCP port number bound with the **bind(2)** system call after the socket is created. By default, TCP sockets are active. A passive socket is created by calling the **listen(2)** system call after binding the socket with **bind**. This establishes a queuing parameter for the passive socket. After this, connections to the passive socket can be received with the **accept(2)** system call. Active sockets use the **connect(2)** call after binding to initiate connections.

By using the special value **INADDR\_ANY**, the local IP address can be left unspecified in the **bind** call by either active or passive TCP sockets. This feature is usually used if the local address is either unknown or irrelevant. If left unspecified, the local IP address will be bound at connection time to the address of the network interface used to service the connection.

Once a connection has been established, data can be exchanged using the **read(2V)** and **write(2V)** system calls.

TCP supports one socket option which is set with **setsockopt** and tested with **getsockopt(2)**. Under most circumstances, TCP sends data when it is presented. When outstanding data has not yet been acknowledged, it gathers small amounts of output to be sent in a single packet once an acknowledgement is received. For a small number of clients, such as window systems that send a stream of mouse events which receive no replies, this packetization may cause significant delays. Therefore, TCP provides a boolean option, **TCP\_NODELAY** (defined in **<netinet/tcp.h>**), to defeat this algorithm. The option level for the **setsockopt** call is the protocol number for TCP, available from **getprotobyname** (see **getprotoent(3N)**).

Options at the IP level may be used with TCP; see **ip(4P)**.

TCP provides an urgent data mechanism, which may be invoked using the out-of-band provisions of **send(2)**. The caller may mark one byte as "urgent" with the **MSG\_OOB** flag to **send(2)**. This causes an "urgent pointer" pointing to this byte to be set in the TCP stream. The receiver on the other side of the stream is notified of the urgent data by a **SIGURG** signal. The **SIOCATMARK** ioctl returns a value indicating whether the stream is at the urgent mark. Because the system never returns data across the urgent mark in a single **read(2V)** call, it is possible to advance to the urgent data in a simple loop which reads data, testing the socket with the **SIOCATMARK** ioctl, until it reaches the mark.

Incoming connection requests that include an IP source route option are noted, and the reverse source route is used in responding.

TCP assumes the datagram service it is layered above is unreliable. A checksum over all data helps TCP implement reliability. Using a window-based flow control mechanism that makes use of positive acknowledgements, sequence numbers, and a retransmission strategy, TCP can usually recover when datagrams are damaged, delayed, duplicated or delivered out of order by the underlying communication medium.

If the local TCP receives no acknowledgements from its peer for a period of time, as would be the case if the remote machine crashed, the connection is closed and an error is returned to the user. If the remote machine reboots or otherwise loses state information about a TCP connection, the connection is aborted and an error is returned to the user.

#### ERRORS

A socket operation may fail if:

EISCONN	A <b>connect</b> operation was attempted on a socket on which a <b>connect</b> operation had already been performed.
ETIMEDOUT	A connection was dropped due to excessive retransmissions.
ECONNRESET	The remote peer forced the connection to be closed (usually because the remote machine has lost state information about the connection due to a crash).
ECONNREFUSED	The remote peer actively refused connection establishment (usually because no process is listening to the port).
EADDRINUSE	A <b>bind</b> operation was attempted on a socket with a network address/port pair that has already been bound to another socket.
EADDRNOTAVAIL	A <b>bind</b> operation was attempted on a socket with a network address for which no network interface exists.
EACCES	A <b>bind</b> operation was attempted with a "reserved" port number and the effective user ID of the process was not super-user.
ENOBUFS	The system ran out of memory for internal data structures.

#### SEE ALSO

**accept(2)**, **bind(2)**, **connect(2)**, **getsockopt(2)**, **listen(2)**, **read(2V)**, **send(2)**, **write(2V)**, **getprotoent(3N)**, **inet(4F)**, **ip(4P)**

Postel, Jon, *Transmission Control Protocol - DARPA Internet Program Protocol Specification*, RFC 793, Network Information Center, SRI International, Menlo Park, Calif., September 1981.

#### BUGS

**SIOCShiwat** and **SIOCGhiwat** `ioctl`'s to set and get the high water mark for the socket queue, and so that it can be changed from 2048 bytes to be larger or smaller, have been defined (in `<sys/ioctl.h>`) but not implemented.

**NAME**

tcptli – TLI-Conforming TCP Stream-Head

**CONFIG**

**pseudo-device clone**

**pseudo-device tcptli32**

**SYNOPSIS**

```
#include <fcntl.h>
#include <netli/tiuser.h>

tfd = t_open("/dev/tcp", O_RDWR, tinfo);
struct t_info *tinfo;
```

**DESCRIPTION**

TCPTLI provides access to TCP service via the Transport Library Interface (TLI). Prior to this release, TCP access was only possible via the socket programming interface. Programmers have the choice of using either the socket or TLI programming interface for their application.

TCPTLI is implemented in STREAMS conforming to the Transport Provider Interface (TPI) specification as a TCP Transport Provider to a TLI application. It utilizes the existing underlying socket and TCP support in the SunOS kernel to communicate over the network. It is also a clone driver, see **clone(4)** for more characteristics pertaining to a clone STREAMS driver.

The notion of an address is the same as the socket address (struct `sockaddr_in`) defined in `<netinet/in.h>`. TCPTLI maintains transport state information for each outstanding connection and the current state of the provider may be retrieved via the `t_getstate(3N)` call. See `t_getstate(3N)` for a list of possible states.

A server usually starts up with the `t_open(3N)` call followed by `t_bind(3N)` to bind an address that it listens for incoming connection. It may call `t_listen(3N)` to retrieve an indication of a connect request from another transport user, and then calls `t_accept(3N)` if it is willing to provide its service. TLI allows a server to accept connection on the same file descriptor it is listening on, or a different file descriptor (as in the sense of socket's `accept(2)`).

A client usually calls `t_open(3N)` and followed by a call to `t_bind(3N)`. Then it calls `t_connect(3N)` to the address of a server advertized for providing service. Once the connection is established, it may use `t_rcv(3N)` and `t_snd(3N)` to receive and send data. The routine `t_close(3N)` is used to terminate the connection.

**TLI ERRORS**

An TLI operation may fail if one of the following error conditions is encountered. They are returned by the TLI user level library.

TBADADDR	Incorrect/invalid address format supplied by the user.
TBADOPT	Incorrect option.
TACCESS	No permission.
TBADF	Illegal transport file descriptor.
TNOADDR	Could not allocate address
TOUTSTATE	The transport is in an incorrect state.
TBADSEQ	Incorrect sequence number.
TSYSERR	A system error, i.e. below the transport level (see list below) is encountered.
TLOOK	An event requires attention.
TBADDATA	Illegal amount of data
TBUFOVFLW	Buffer not large enough.

TFLOW	Flow control problem.
TNODATA	No data.
TNODIS	No <code>discon_ind</code> is found on the queue.
TNOUDERR	Unit data not found.
TBADFLAG	Bad flags.
TNOREL	No orderly release request found on queue.
TNOTSUPPORT	Protocol/primitive is not supported.
TSTATECHNG	State is in the process of changing.

**SYSTEM ERRORS**

The following errors are returned by TCPTLI. However they may be translated to the above TLI errors by the user level library ( `libnsl` ).

ENXIO	Invalid device or address, out of range.
EBUSY	Request device is busy or not ready.
ENOMEM	Not enough memory for transmitting data, non fatal.
EPROTO	The operation encountered an underlying protocol. error (TCP).
EWOULDBLOCK	The operation would block as normally the file descriptors are set with non-blocking flag.
EACCES	Permission denied.
ENOBUFS	The system ran out of memory for internal (network) data structures.

**SEE ALSO**

`accept(2)`, `t_open(3N)`, `t_close(3N)`, `t_accept(3N)`, `t_getstate(3N)`, `t_bind(3N)`, `t_connect(3N)`, `t_rcv(3N)`, `t_snd(3N)`, `t_alloc(3N)`, `t_unbind(3N)`, `t_getinfo(3N)`

**BUGS**

Only TCP (i.e. connection oriented) protocol is supported, no UDP. The maximum network connection is 32 by default. A new kernel has to be configured if an increase of such limit is desired: by changing the entry `pseudo-device tcptli32` in the kernel config file to `tcptli64`.

**NAME**

termio – general terminal interface

**SYNOPSIS**

```
#include <sys/termios.h>
```

**DESCRIPTION**

Asynchronous communications ports, pseudo-terminals, and the special interface accessed by `/dev/tty` all use the same general interface, no matter what hardware (if any) is involved. The remainder of this section discusses the common features of this interface.

**Opening a Terminal Device File**

When a terminal file is opened, the process normally waits until a connection is established. In practice, users' programs seldom open these files; they are opened by `getty(8)` and become a user's standard input, output, and error files. The state of the software carrier flag will effect the ability to open a line.

**Sessions**

Processes are now grouped by session, then process group, then process id. Each session is associated with one "login" session (windows count as logins). A process creates a session by calling `setsid(2V)`, which will put the process in a new session as its only member and as the session leader of that session.

**Process Groups**

A terminal may have a distinguished process group associated with it. This distinguished process group plays a special role in handling signal-generating input characters, as discussed below in the **Special Characters** section below. The terminal's process group can be set only to process groups that are members of the terminal's session.

A command interpreter, such as `csh(1)`, that supports "job control" can allocate the terminal to different *jobs*, or process groups, by placing related processes in a single process group and associating this process group with the terminal. A terminal's associated process group may be set or examined by a process with sufficient privileges. The terminal interface aids in this allocation by restricting access to the terminal by processes that are not in the current process group; see **Job Access Control** below.

**Orphaned Process Groups**

An orphaned process group is a process group that has no parent, in a different process group, and in the same session. In other words, there is no process that can handle job control signals for the process group.

**The Controlling Terminal**

A terminal may belong to a process as its *controlling terminal*. If a process that is a session leader, and that does not have a controlling terminal, opens a terminal file not already associated with a session, the terminal associated with that terminal file becomes the controlling terminal for that process, and the terminal's distinguished process group is set to the process group of that process. (Currently, this also happens if a process that does not have a controlling terminal and is not a member of a process group opens a terminal. In this case, if the terminal is not associated with a session, a new session is created with a process group ID equal to the process ID of the process in question, and the terminal is assigned to that session. The process is made a member of the terminal's process group.)

If a process does not wish to acquire the terminal as a controlling terminal (as is the case with many daemons that open `/dev/console`), the process should or `O_NOCTTY` into the second argument to `open(2V)`.

The controlling terminal is inherited by a child process during a `fork(2V)`. A process relinquishes its control terminal when it changes its process group using `setsid(2V)`, when it tries to change back to process group 0 via a `setpgrp(2V)` with arguments (`mygid, 0`), or when it issues a `TIOCNOTTY ioctl(2)` call on a file descriptor created by opening the file `/dev/tty`. Both of the last two cases cause a `setsid(2V)` to be called on the process' behalf. This is an attempt to allow old binaries (that couldn't have known about `setsid(2V)`) to still acquire controlling terminals. It doesn't always work, see `setsid(8V)` for a workaround for those cases.



When a session leader that has a controlling terminal terminates, the distinguished process group of the controlling terminal is set to zero (indicating no distinguished process group). This allows the terminal to be acquired as a controlling terminal by a new session leader.

#### Closing a Terminal Device File

When a terminal device file is closed, the process closing the file waits until all output is drained; all pending input is then flushed, and finally a disconnect is performed. If HUPCL is set, the existing connection is severed (by hanging up the phone line, if appropriate).

#### Job Access Control

If a process is in the (non-zero) distinguished process group of its controlling terminal (if this is true, the process is said to be a *foreground process*), then `read(2V)` operations are allowed as described below in **Input Processing and Reading Characters**. If a process is not in the (non-zero) distinguished process group of its controlling terminal (if this is true, the process is said to be a *background process*), then any attempts to read from that terminal will typically send that process' process group a SIGTTIN signal. If the process is ignoring SIGTTIN, has SIGTTIN blocked, is a member of an orphaned process group, or is in the middle of process creation using `vfork(2)`, the read will return `-1` and set `errno` to `EIO`, and the SIGTTIN signal will not be sent. The SIGTTIN signal will normally stop the members of that process group.

When the TOSTOP bit is set in the `c_lflag` field, attempts by a background process to write to its controlling terminal will typically send that process' process group a SIGTTOU signal. If the process is ignoring SIGTTOU, has SIGTTOU blocked, or is in the middle of process creation using `vfork()`, the process will be allowed to write to the terminal and the SIGTTOU signal will not be sent. If the process is orphaned, the write will return `-1` and set `errno` to `EIO`, and the SIGTTOU signal will not be sent. SIGTTOU signal will normally stop the members of that process group. Certain `ioctl()` calls that set terminal parameters are treated in this same fashion, except that TOSTOP is not checked; the effect is identical to that of terminal writes when TOSTOP is set. See **IOCTLS**.

#### Input Processing and Reading Characters

A terminal associated with one of these files ordinarily operates in full-duplex mode. Characters may be typed at any time, even while output is occurring, and are only lost when the system's character input buffers become completely full, which is rare, or when the user has accumulated the maximum allowed number of input characters that have not yet been read by some program. This limit is available is `{MAX_CANON}` characters (see `pathconf(2V)`). If the IMAXBEL mode has not been selected, all the saved characters are thrown away without notice when the input limit is reached; if the IMAXBEL mode has been selected, the driver refuses to accept any further input, and echoes a bell (ASCII BEL).

Two general kinds of input processing are available, determined by whether the terminal device file is in canonical mode or non-canonical mode (see **ICANON** in the **Local Modes** section).

The style of input processing can also be very different when the terminal is put in non-blocking I/O mode; see `read(2V)`. In this case, reads from the terminal will never block.

It is possible to simulate terminal input using the `TIOCSTI ioctl()` call, which takes, as its third argument, the address of a character. The system pretends that this character was typed on the argument terminal, which must be the process' controlling terminal unless the process' effective user ID is super-user.

#### Canonical Mode Input Processing

In canonical mode input processing, terminal input is processed in units of lines. A line is delimited by a NEWLINE (ASCII LF) character, an EOF (by default, an ASCII EOT) character, or one of two user-specified end-of-line characters, `EOL` and `EOL2`. This means that a `read()` will not complete until an entire line has been typed or a signal has been received. Also, no matter how many characters are requested in the read call, at most one line will be returned. It is not, however, necessary to read a whole line at once; any number of characters may be requested in a read, even one, without losing information.

Erase and kill processing occurs during input. The ERASE character (by default, the character DEL) erases the last character typed in the current input line. The WERASE character (by default, the character CTRL-W) erases the last "word" typed in the current input line (but not any preceding SPACE or TAB characters). A "word" is defined as a sequence of non-blank characters, with TAB characters counted as blanks.

Neither **ERASE** nor **WERASE** will erase beyond the beginning of the line. The **KILL** character (by default, the character **CTRL-U**) kills (deletes) the entire current input line, and optionally outputs a **NEWLINE** character. All these characters operate on a key-stroke basis, independently of any backspacing or tabbing that may have been done.

The **REPRINT** character (the character **CTRL-R**) prints a **NEWLINE** followed by all characters that have not been read. Reprinting also occurs automatically if characters that would normally be erased from the screen are fouled by program output. The characters are reprinted as if they were being echoed; as a consequence, if **ECHO** is not set, they are not printed.

The **ERASE** and **KILL** characters may be entered literally by preceding them with the escape character (**\**). In this case the escape character is not read. The **ERASE** and **KILL** characters may be changed.

#### Non-Canonical Mode Input Processing

In non-canonical mode input processing, input characters are not assembled into lines, and erase and kill processing does not occur. The **MIN** and **TIME** values are used to determine how to process the characters received.

**MIN** represents the minimum number of characters that should be received when the read is satisfied (when the characters are returned to the user). **TIME** is a timer of 0.10 second granularity that is used to timeout bursty and short term data transmissions. The four possible values for **MIN** and **TIME** and their interactions are described below.

#### Case A: **MIN > 0, TIME > 0**

In this case **TIME** serves as an intercharacter timer and is activated after the first character is received. Since it is an intercharacter timer, it is reset after a character is received. The interaction between **MIN** and **TIME** is as follows: as soon as one character is received, the intercharacter timer is started. If **MIN** characters are received before the intercharacter timer expires (remember that the timer is reset upon receipt of each character), the read is satisfied. If the timer expires before **MIN** characters are received, the characters received to that point are returned to the user. Note: if **MIN** expires at least one character will be returned because the timer would not have been enabled unless a character was received. In this case (**MIN > 0, TIME > 0**) the read will sleep until the **MIN** and **TIME** mechanisms are activated by the receipt of the first character.

#### Case B: **MIN > 0, TIME = 0**

In this case, since the value of **TIME** is zero, the timer plays no role and only **MIN** is significant. A pending read is not satisfied until **MIN** characters are received (the pending read will sleep until **MIN** characters are received). A program that uses this case to read record-based terminal I/O may block indefinitely in the read operation.

#### Case C: **MIN = 0, TIME > 0**

In this case, since **MIN = 0**, **TIME** no longer represents an intercharacter timer. It now serves as a read timer that is activated as soon as a **read()** is done. A read is satisfied as soon as a single character is received or the read timer expires. Note: in this case if the timer expires, no character will be returned. If the timer does not expire, the only way the read can be satisfied is if a character is received. In this case the read will not block indefinitely waiting for a character – if no character is received within **TIME\*.10** seconds after the read is initiated, the read will return with zero characters.

#### Case D: **MIN = 0, TIME = 0**

In this case return is immediate. The minimum of either the number of characters requested or the number of characters currently available will be returned without waiting for more characters to be input.

#### Comparison of the Different Cases of **MIN, TIME** Interaction

Some points to note about **MIN** and **TIME**:

- In the following explanations one may notice that the interactions of **MIN** and **TIME** are not symmetric. For example, when **MIN > 0** and **TIME = 0**, **TIME** has no effect. However, in the opposite case where **MIN = 0** and **TIME > 0**, both **MIN** and **TIME** play a role in that **MIN** is satisfied with the receipt of a single character.

- Also note that in case A ( $\text{MIN} > 0$ ,  $\text{TIME} > 0$ ),  $\text{TIME}$  represents an intercharacter timer while in case C ( $\text{TIME} = 0$ ,  $\text{TIME} > 0$ )  $\text{TIME}$  represents a read timer.

These two points highlight the dual purpose of the  $\text{MIN}/\text{TIME}$  feature. Cases A and B, where  $\text{MIN} > 0$ , exist to handle burst mode activity (for example, file transfer programs) where a program would like to process at least  $\text{MIN}$  characters at a time. In case A, the intercharacter timer is activated by a user as a safety measure; while in case B, it is turned off.

Cases C and D exist to handle single character timed transfers. These cases are readily adaptable to screen-based applications that need to know if a character is present in the input queue before refreshing the screen. In case C the read is timed; while in case D, it is not.

Another important note is that  $\text{MIN}$  is always just a minimum. It does not denote a record length. That is, if a program does a read of 20 bytes,  $\text{MIN}$  is 10, and 25 characters are present, 20 characters will be returned to the user.

#### Writing Characters

When one or more characters are written, they are transmitted to the terminal as soon as previously-written characters have finished typing. Input characters are echoed as they are typed if echoing has been enabled. If a process produces characters more rapidly than they can be typed, it will be suspended when its output queue exceeds some limit. When the queue has drained down to some threshold, the program is resumed.

#### Special Characters

Certain characters have special functions on input and/or output. These functions and their default character values are summarized as follows:

<b>INTR</b>	(CTRL-C or ASCII ETX) generates a SIGINT signal, which is sent to all processes in the distinguished process group associated with the terminal. Normally, each such process is forced to terminate, but arrangements may be made either to ignore the signal or to receive a trap to an agreed-upon location; see <code>sigvec(2)</code> .
<b>QUIT</b>	(CTRL-  or ASCII FS) generates a SIGQUIT signal, which is sent to all processes in the distinguished process group associated with the terminal. Its treatment is identical to the interrupt signal except that, unless a receiving process has made other arrangements, it will not only be terminated but a core image file (called <code>core</code> ) will be created in the current working directory.
<b>ERASE</b>	(Rubout or ASCII DEL) erases the preceding character. It will not erase beyond the start of a line, as delimited by a NL, EOF, EOL, or EOL2 character.
<b>WERASE</b>	(CTRL-W or ASCII ETB) erases the preceding "word". It will not erase beyond the start of a line, as delimited by a NL, EOF, EOL, or EOL2 character.
<b>KILL</b>	(CTRL-U or ASCII NAK) deletes the entire line, as delimited by a NL, EOF, EOL, or EOL2 character.
<b>REPRINT</b>	(CTRL-R or ASCII DC2) reprints all characters that have not been read, preceded by a NEWLINE.
<b>EOF</b>	(CTRL-D or ASCII EOT) may be used to generate an end-of-file from a terminal. When received, all the characters waiting to be read are immediately passed to the program, without waiting for a NEWLINE, and the EOF is discarded. Thus, if there are no characters waiting, which is to say the EOF occurred at the beginning of a line, zero characters will be passed back, which is the standard end-of-file indication.
<b>NL</b>	(ASCII LF) is the normal line delimiter. It can not be changed; it can, however, be escaped by the LNEXT character.
<b>EOL</b>	
<b>EOL2</b>	(ASCII NUL) are additional line delimiters, like NL. They are not normally used.

- SUSP** (CTRL-Z or ASCII EM) is used by the job control facility to change the current job to return to the controlling job. It generates a SIGTSTP signal, which stops all processes in the terminal's process group.
- STOP** (CTRL-S or ASCII DC3) can be used to temporarily suspend output. It is useful with CRT terminals to prevent output from disappearing before it can be read. While output is suspended, STOP characters are ignored and not read.
- START** (CTRL-Q or ASCII DC1) is used to resume output that has been suspended by a STOP character. While output is not suspended, START characters are ignored and not read.
- DISCARD** (CTRL-O or ASCII SI) causes subsequent output to be discarded until another DISCARD character is typed, more input arrives, or the condition is cleared by a program.
- LNEXT** (CTRL-V or ASCII SYN) causes the special meaning of the next character to be ignored; this works for all the special characters mentioned above. This allows characters to be input that would otherwise get interpreted by the system (for example, KILL, QUIT.)

The character values for INTR, QUIT, ERASE, WERASE, KILL, REPRINT, EOF, EOL, EOL2, SUSP, STOP, START, DISCARD, and LNEXT may be changed to suit individual tastes. If the value of a special control character is 0, the function of that special control character will be disabled. The ERASE, KILL, and EOF characters may be escaped by a preceding \ character, in which case no special function is done. Any of the special characters may be preceded by the LNEXT character, in which case no special function is done.

If IEXTEN is added to the local modes (this is the default), then all of the special characters are in effect. If IEXTEN is cleared from the local modes, then only the following POSIX.1 compatible specials are seen as specials: INTR, QUIT, ERASE, KILL, EOF, NL, EOL, SUSP, STOP, START, and CR.

#### Software Carrier Mode

The software carrier mode can be enabled or disabled using the TIOCSSOFTCAR `ioctl()`. If the software carrier flag for a line is off, the line pays attention to the hardware carrier detect (DCD) signal. The `tty` device associated with the line can not be opened until DCD is asserted. If the software carrier flag is on, the line behaves as if DCD is always asserted.

The software carrier flag is usually turned on for locally connected terminals or other devices, and is off for lines with modems.

To be able to issue the TIOCGSOFTCAR and TIOCSSOFTCAR `ioctl()` calls, the `tty` line should be opened with `O_NDELAY` so that the `open(2V)` will not wait for the carrier.

#### Modem Disconnect

If a modem disconnect is detected, and the CLOCAL flag is not set in the `c_cflag` field, a SIGHUP signal is sent to all processes in the distinguished process group associated with this terminal. Unless other arrangements have been made, this signal terminates the processes. If SIGHUP is ignored or caught, any subsequent `read()` returns with an end-of-file indication until the terminal is closed. Thus, programs that read a terminal and test for end-of-file can terminate appropriately after a disconnect. Any subsequent `write()` will return `-1` and set `errno` to `EIO` until the terminal is closed.

A SIGHUP signal is sent to the `tty` if the software carrier flag is off and the hardware carrier detect drops.

#### Terminal Parameters

The parameters that control the behavior of devices and modules providing the `termios` interface are specified by the `termios` structure, defined by `<sys/termios.h>`. Several `ioctl()` system calls that fetch or change these parameters use this structure:

```
#define NCCS      17
struct termios {
    unsigned long  c_iflag;    /* input modes */
    unsigned long  c_oflag;    /* output modes */
    unsigned long  c_cflag;    /* control modes */
```

```

        unsigned long   c_iflag;    /* local modes */
        unsigned char  c_line;     /* line discipline */
        unsigned char  c_cc[NCCS]; /* control chars */
};

```

The special control characters are defined by the array `c_cc`. The relative positions and initial values for each function are as follows:

0	VINTR	ETX
1	VQUIT	FS
2	VERASE	DEL
3	VKILL	NAK
4	VEOF	EOT
5	VEOL	NUL
6	VEOL2	NUL
7	VSWTCH	NUL
8	VSTART	DC1
9	VSTOP	DC3
10	VSUSP	EM
12	VREPRINT	DC2
13	VDISCARD	SI
14	VWERASE	ETB
15	VLNEXT	SYN

The `MIN` value is stored in the `VMIN` element of the `c_cc` array, and the `TIME` value is stored in the `VTIME` element of the `c_cc` array. The `VMIN` element is the same element as the `VEOF` element, and the `VTIME` element is the same element as the `VEOL` element.

#### Input Modes

The `c_iflag` field describes the basic terminal input control:

<code>IGNBRK</code>	0000001	Ignore break condition.
<code>BRKINT</code>	0000002	Signal interrupt on break.
<code>IGNPAR</code>	0000004	Ignore characters with parity errors.
<code>PARMRK</code>	0000010	Mark parity errors.
<code>INPCK</code>	0000020	Enable input parity check.
<code>ISTRIP</code>	0000040	Strip character.
<code>INLCR</code>	0000100	Map NL to CR on input.
<code>IGNCR</code>	0000200	Ignore CR.
<code>ICRNL</code>	0000400	Map CR to NL on input.
<code>IUCLC</code>	0001000	Map upper-case to lower-case on input.
<code>IXON</code>	0002000	Enable start/stop output control.
<code>IXANY</code>	0004000	Enable any character to restart output.
<code>IXOFF</code>	0010000	Enable start/stop input control.
<code>IMAXBEL</code>	0020000	Echo BEL on input line too long.

If `IGNBRK` is set, a break condition (a character framing error with data all zeros) detected on input is ignored, that is, not put on the input queue and therefore not read by any process. Otherwise, if `BRKINT` is set, a break condition will generate a `SIGINT` and flush both the input and output queues. If neither `IGNBRK` nor `BRKINT` is set, a break condition is read as a single ASCII NUL character (`^0`).

If `IGNPAR` is set, characters with framing or parity errors (other than break) are ignored. Otherwise, if `PARMRK` is set, a character with a framing or parity error that is not ignored is read as the three-character sequence: `^377`, `^0`, `X`, where `X` is the data of the character received in error. To avoid ambiguity in this case, if `ISTRIP` is not set, a valid character of `^377` is read as `^377`, `^377`. If neither `IGNPAR` nor `PARMRK` is set, a framing or parity error (other than break) is read as a single ASCII NUL character (`^0`).

If **INPCK** is set, input parity checking is enabled. If **INPCK** is not set, input parity checking is disabled. This allows output parity generation without input parity errors.

If **ISTRIP** is set, valid input characters are first stripped to 7 bits, otherwise all 8 bits are processed.

If **INLCR** is set, a received **NL** character is translated into a **CR** character. If **IGNCR** is set, a received **CR** character is ignored (not read). Otherwise if **ICRNL** is set, a received **CR** character is translated into a **NL** character.

If **IUCLC** is set, a received upper-case alphabetic character is translated into the corresponding lower-case character.

If **IXON** is set, start/stop output control is enabled. A received **STOP** character will suspend output and a received **START** character will restart output. The **STOP** and **START** characters will not be read, but will merely perform flow control functions. If **IXANY** is set, any input character will restart output that has been suspended.

If **IXOFF** is set, the system will transmit a **STOP** character when the input queue is nearly full, and a **START** character when enough input has been read that the input queue is nearly empty again.

If **IMAXBEL** is set, the ASCII **BEL** character is echoed if the input stream overflows. Further input will not be stored, but any input already present in the input stream will not be disturbed. If **IMAXBEL** is not set, no **BEL** character is echoed, and all input present in the input queue is discarded if the input stream overflows.

The initial input control value is **BRKINT**, **ICRNL**, **IXON**, **ISTRIP**.

#### Output modes

The **c\_oflag** field specifies the system treatment of output:

<b>OPOST</b>	0000001	Postprocess output.
<b>OLCUC</b>	0000002	Map lower case to upper on output.
<b>ONLCR</b>	0000004	Map <b>NL</b> to <b>CR-NL</b> on output.
<b>OCRNL</b>	0000010	Map <b>CR</b> to <b>NL</b> on output.
<b>ONOCR</b>	0000020	No <b>CR</b> output at column 0.
<b>ONLRET</b>	0000040	<b>NL</b> performs <b>CR</b> function.
<b>OFILL</b>	0000100	Use fill characters for delay.
<b>OFDEL</b>	0000200	Fill is <b>DEL</b> , else <b>NUL</b> .
<b>NLDLY</b>	0000400	Select new-line delays:
<b>NL0</b>	0	
<b>NL1</b>	0000400	
<b>CRDLY</b>	0003000	Select carriage-return delays:
<b>CR0</b>	0	
<b>CR1</b>	0001000	
<b>CR2</b>	0002000	
<b>CR3</b>	0003000	
<b>TABDLY</b>	0014000	Select horizontal-tab delays:
<b>TAB0</b>	0	or tab expansion:
<b>TAB1</b>	0004000	
<b>TAB2</b>	0010000	
<b>XTABS</b>	0014000	Expand tabs to spaces.
<b>BSDLY</b>	0020000	Select backspace delays:
<b>BS0</b>	0	
<b>BS1</b>	0020000	
<b>VTDLY</b>	0040000	Select vertical-tab delays:
<b>VT0</b>	0	
<b>VT1</b>	0040000	

```

FFDLY  0100000 Select form-feed delays:
FF0    0
FF1    0100000

```

If **OPOST** is set, output characters are post-processed as indicated by the remaining flags, otherwise characters are transmitted without change.

If **OLCUC** is set, a lower-case alphabetic character is transmitted as the corresponding upper-case character. This function is often used in conjunction with **IUCLC**.

If **ONLCR** is set, the **NL** character is transmitted as the **CR-NL** character pair. If **OCRNL** is set, the **CR** character is transmitted as the **NL** character. If **ONOCR** is set, no **CR** character is transmitted when at column 0 (first position). If **ONLRET** is set, the **NL** character is assumed to do the carriage-return function; the column pointer will be set to 0 and the delays specified for **CR** will be used. Otherwise the **NL** character is assumed to do just the line-feed function; the column pointer will remain unchanged. The column pointer is also set to 0 if the **CR** character is actually transmitted.

The delay bits specify how long transmission stops to allow for mechanical or other movement when certain characters are sent to the terminal. In all cases a value of 0 indicates no delay. If **OFILL** is set, fill characters will be transmitted for delay instead of a timed delay. This is useful for high baud rate terminals that need only a minimal delay. If **OFDEL** is set, the fill character is **DEL**, otherwise **NUL**.

If a form-feed or vertical-tab delay is specified, it lasts for about 2 seconds.

New-line delay lasts about 0.10 seconds. If **ONLRET** is set, the **RETURN** delays are used instead of the **NEWLINE** delays. If **OFILL** is set, two fill characters will be transmitted.

Carriage-return delay type 1 is dependent on the current column position, type 2 is about 0.10 seconds, and type 3 is about 0.15 seconds. If **OFILL** is set, delay type 1 transmits two fill characters, and type 2, four fill characters.

Horizontal-tab delay type 1 is dependent on the current column position. Type 2 is about 0.10 seconds. Type 3, specified by **TAB3** or **XTABS**, specifies that **TAB** characters are to be expanded into **SPACE** characters. If **OFILL** is set, two fill characters will be transmitted for any delay.

Backspace delay lasts about 0.05 seconds. If **OFILL** is set, one fill character will be transmitted.

The actual delays depend on line speed and system load.

The initial output control value is **OPOST**, **ONLCR**, **XTABS**.

The **c\_cflag** field describes the hardware control of the terminal:

<b>CBAUD</b>	0000017	Baud rate:
<b>B0</b>	0	Hang up
<b>B50</b>	0000001	50 baud
<b>B75</b>	0000002	75 baud
<b>B110</b>	0000003	110 baud
<b>B134</b>	0000004	134.5 baud
<b>B150</b>	0000005	150 baud
<b>B200</b>	0000006	200 baud
<b>B300</b>	0000007	300 baud
<b>B600</b>	0000010	600 baud
<b>B1200</b>	0000011	1200 baud
<b>B1800</b>	0000012	1800 baud
<b>B2400</b>	0000013	2400 baud
<b>B4800</b>	0000014	4800 baud
<b>B9600</b>	0000015	9600 baud
<b>B19200</b>	0000016	19200 baud
<b>B38400</b>	0000017	38400 baud

<b>CSIZE</b>	0000060	Character size:
<b>CS5</b>	0	5 bits
<b>CS6</b>	0000020	6 bits
<b>CS7</b>	0000040	7 bits
<b>CS8</b>	0000060	8 bits
<b>CSTOPB</b>	0000100	Send two stop bits, else one.
<b>CREAD</b>	0000200	Enable receiver.
<b>PARENB</b>	0000400	Parity enable.
<b>PARODD</b>	0001000	Odd parity, else even.
<b>HUPCL</b>	0002000	Hang up on last close.
<b>CLOCAL</b>	0004000	Local line, else dial-up.
<b>CIBAUD</b>	03600000	Input baud rate, if different from output rate.
<b>CRTSCTS</b>	020000000000	Enable RTS/CTS flow control.

The **CBAUD** bits specify the baud rate. The zero baud rate, **B0**, is used to hang up the connection. If **B0** is specified, the modem control lines will cease to be asserted. Normally, this will disconnect the line. If the **CIBAUD** bits are not zero, they specify the input baud rate, with the **CBAUD** bits specifying the output baud rate; otherwise, the output and input baud rates are both specified by the **CBAUD** bits. The values for the **CIBAUD** bits are the same as the values for the **CBAUD** bits, shifted left **IBSHIFT** bits. For any particular hardware, impossible speed changes are ignored.

The **CSIZE** bits specify the character size in bits for both transmission and reception. This size does not include the parity bit, if any. If **CSTOPB** is set, two stop bits are used, otherwise one stop bit. For example, at 110 baud, two stop bits are required.

If **PARENB** is set, parity generation and detection is enabled and a parity bit is added to each character. If parity is enabled, the **PARODD** flag specifies odd parity if set, otherwise even parity is used.

If **CREAD** is set, the receiver is enabled. Otherwise no characters will be received.

If **HUPCL** is set, the modem control lines for the port will be disconnected when the last process with the line open closes it or terminates.

If **CLOCAL** is set, a connection does not depend on the state of the modem status lines. Otherwise modem control is assumed.

If **CRTSCTS** is set, and the terminal has modem control lines associated with it, the Request To Send (RTS) modem control line will be raised, and output will occur only if the Clear To Send (CTS) modem status line is raised. If the CTS modem status line is lowered, output is suspended until CTS is raised. Some hardware may not support this function, and other hardware may not permit it to be disabled; in either of these cases, the state of the **CRTSCTS** flag is ignored.

The initial hardware control value after open is **B9600**, **CS7**, **CREAD**, **PARENB**.

#### Local Modes

The **c\_iflag** field of the argument structure is used by the line discipline to control terminal functions. The basic line discipline provides the following:

<b>ISIG</b>	0000001	Enable signals.
<b>ICANON</b>	0000002	Canonical input (erase and kill processing).
<b>XCASE</b>	0000004	Canonical upper/lower presentation.
<b>ECHO</b>	0000010	Enable echo.
<b>ECHOE</b>	0000020	Echo erase character as BS-SP-BS.
<b>ECHOK</b>	0000040	Echo NL after kill character.
<b>ECHONL</b>	0000100	Echo NL.
<b>NOFLSH</b>	0000200	Disable flush after interrupt or quit.
<b>TOSTOP</b>	0000400	Send SIGTTOU for background output.
<b>ECHOCTL</b>	0001000	Echo control characters as <i>^char</i> , delete as <i>^?</i> .
<b>ECHOPRT</b>	0002000	Echo erase character as character erased.
<b>ECHOKE</b>	0004000	BS-SP-BS erase entire line on line kill.



FLUSHO	0020000	Output is being flushed.
PENDIN	0040000	Retype pending input at next read or input character.
IEXTEN	0100000	Recognize all specials (if clear, POSIX only).

If **ISIG** is set, each input character is checked against the special control characters **INTR**, **QUIT**, and **SUSP**. If an input character matches one of these control characters, the function associated with that character is performed. If **ISIG** is not set, no checking is done. Thus these special input functions are possible only if **ISIG** is set.

If **ICANON** is set, canonical processing is enabled. This is affected by the **IEXTEN** bit (see **Special Characters** above). This enables the erase, word erase, kill, and reprint edit functions, and the assembly of input characters into lines delimited by **NL**, **EOF**, **EOL**, and **EOL2**. If **ICANON** is not set, read requests are satisfied directly from the input queue. A read will not be satisfied until at least **MIN** characters have been received or the timeout value **TIME** has expired between characters. This allows fast bursts of input to be read efficiently while still allowing single character input. The time value represents tenths of seconds. See the *Non-canonical Mode Input Processing* section for more details.

If **XCASE** is set, and if **ICANON** is set, an upper-case letter is accepted on input by preceding it with a **\** character, and is output preceded by a **\** character. In this mode, the following escape sequences are generated on output and accepted on input:

for:	use:
<b>\</b>	<b>\</b>
<b> </b>	<b>\!</b>
<b>-</b>	<b>\^</b>
<b>{</b>	<b>\(</b>
<b>}</b>	<b>\)</b>
<b>\</b>	<b>\\</b>

For example, **A** is input as **\a**, **\n** as **\\n**, and **\N** as **\\N**.

If **ECHO** is set, characters are echoed as received. If **ECHO** is not set, input characters are not echoed.

If **ECHOCTL** is not set, all control characters (characters with codes between 0 and 37 octal) are echoed as themselves. If **ECHOCTL** is set, all control characters other than ASCII **TAB**, ASCII **NL**, the **START** character, and the **STOP** character, are echoed as **^X**, where **X** is the character given by adding 100 octal to the control character's code (so that the character with octal code 1 is echoed as **^A**), and the ASCII **DEL** character, with code 177 octal, is echoed as **^?**.

When **ICANON** is set, the following echo functions are possible:

- If **ECHO** and **ECHOE** are set, and **ECHOPRT** is not set, the **ERASE** and **WERASE** characters are echoed as one or more ASCII **BS SP BS**, which will clear the last character(s) from a CRT screen.
- If **ECHO** and **ECHOPRT** are set, the first **ERASE** and **WERASE** character in a sequence echoes as a backslash (**\**) followed by the characters being erased. Subsequent **ERASE** and **WERASE** characters echo the characters being erased, in reverse order. The next non-erase character types a slash (**/**) before it is echoed.
- If **ECHOKE** is set, the kill character is echoed by erasing each character on the line from the screen (using the mechanism selected by **ECHOE** and **ECHOPRT**).
- If **ECHOK** is set, and **ECHOKE** is not set, the **NL** character will be echoed after the kill character to emphasize that the line will be deleted. Note: an escape character (**\**) or an **LNEXT** character preceding the erase or kill character removes any special function.
- If **ECHONL** is set, the **NL** character will be echoed even if **ECHO** is not set. This is useful for terminals set to local echo (so-called half duplex).

- If ECHOCTL is not set, the EOF character is not echoed, unless it is escaped. Because EOT is the default EOF character, this prevents terminals that respond to EOT from hanging up. If ECHOCTL is set, the EOF character is echoed; if it is not escaped, after it is echoed, one backspace character is output if it is echoed as itself, and two backspace characters are echoed if it is echoed as ^X.

If NOFLSH is set, the normal flush of the input and output queues associated with the INTR, QUIT, and SUSP characters will not be done.

If TOSTOP is set, the signal SIGTTOU is sent to a process that tries to write to its controlling terminal if it is not in the distinguished process group for that terminal. This signal normally stops the process. Otherwise, the output generated by that process is output to the current output stream. Processes that are blocking or ignoring SIGTTOU signals are excepted and allowed to produce output.

If FLUSHO is set, data written to the terminal will be discarded. This bit is set when the FLUSH character is typed. A program can cancel the effect of typing the FLUSH character by clearing FLUSHO.

If PENDIN is set, any input that has not yet been read will be reprinted when the next character arrives as input.

The initial line-discipline control value is ISIG, ICANON, ECHO.

#### Minimum and Timeout

The MIN and TIME values are described above under **Non-canonical Mode Input Processing**. The initial value of MIN is 1, and the initial value of TIME is 0.

#### Termio Structure

The System V **termio** structure is used by other **ioctl()** calls; it is defined by **<sys/termio.h>** as:

```
#define NCC      8
struct termio {
    unsigned short c_iflag;    /* input modes */
    unsigned short c_oflag;    /* output modes */
    unsigned short c_cflag;    /* control modes */
    unsigned short c_lflag;    /* local modes */
    char           c_line;      /* line discipline */
    unsigned char  c_cc[NCC];  /* control chars */
};
```

The special control characters are defined by the array **c\_cc**. The relative positions for each function are as follows:

```
0  VINTR
1  VQUIT
2  VERASE
3  VKILL
4  VEOF
5  VEOL
6  VEOL2
7  reserved
```

The calls that use the **termio** structure only affect the flags and control characters that can be stored in the **termio** structure; all other flags and control characters are unaffected.

#### Terminal Size

The number of lines and columns on the terminal's display (or page, in the case of printing terminals) is specified in the **winsize** structure, defined by **<sys/termios.h>**. Several **ioctl()** system calls that fetch or change these parameters use this structure:

```
struct winsize {
    unsigned short ws_row;    /* rows, in characters */
    unsigned short ws_col;    /* columns, in characters */
};
```

```

        unsigned short    ws_xpixel; /* horizontal size, pixels - not used */
        unsigned short    ws_ypixel; /* vertical size, pixels - not used */
    };

```

#### Modem Lines

On special files representing serial ports, the modem control lines supported by the hardware can be read and the modem status lines supported by the hardware can be changed. The following modem control and status lines may be supported by a device; they are defined by `<sys/termios.h>`:

<code>TIOCM_LE</code>	0001	line enable
<code>TIOCM_DTR</code>	0002	data terminal ready
<code>TIOCM_RTS</code>	0004	request to send
<code>TIOCM_ST</code>	0010	secondary transmit
<code>TIOCM_SR</code>	0020	secondary receive
<code>TIOCM_CTS</code>	0040	clear to send
<code>TIOCM_CAR</code>	0100	carrier detect
<code>TIOCM_RNG</code>	0200	ring
<code>TIOCM_DSR</code>	0400	data set ready

`TIOCM_CD` is a synonym for `TIOCM_CAR`, and `TIOCM_RI` is a synonym for `TIOCM_RNG`.

Not all of these will necessarily be supported by any particular device; check the manual page for the device in question.

#### IOCTLS

The `ioctl()` calls supported by devices and STREAMS modules providing the `termios` interface are listed below. Some calls may not be supported by all devices or modules.

Unless otherwise noted for a specific `ioctl()` call, these functions are restricted from use by background processes. Attempts to perform these calls will cause the process group of the process performing the call to be sent a `SIGTTOU` signal. If the process is ignoring `SIGTTOU`, has `SIGTTOU` blocked, or is in the middle of process creation using `vfork()`, the process will be allowed to perform the call and the `SIGTTOU` signal will not be sent.

<code>TCGETS</code>	The argument is a pointer to a <code>termios</code> structure. The current terminal parameters are fetched and stored into that structure. This call is allowed from a background process; however, the information may subsequently be changed by a foreground process.
<code>TCSETS</code>	The argument is a pointer to a <code>termios</code> structure. The current terminal parameters are set from the values stored in that structure. The change is immediate.
<code>TCSETSW</code>	The argument is a pointer to a <code>termios</code> structure. The current terminal parameters are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted. This form should be used when changing parameters that will affect output.
<code>TCSETSF</code>	The argument is a pointer to a <code>termios</code> structure. The current terminal parameters are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted; all characters queued for input are discarded and then the change occurs.
<code>TCGETA</code>	The argument is a pointer to a <code>termio</code> structure. The current terminal parameters are fetched, and those parameters that can be stored in a <code>termio</code> structure are stored into that structure. This call is allowed from a background process; however, the information may subsequently be changed by a foreground process.
<code>TCSETA</code>	The argument is a pointer to a <code>termio</code> structure. Those terminal parameters that can be stored in a <code>termio</code> structure are set from the values stored in that structure. The change is immediate.

<b>TCSETAW</b>	The argument is a pointer to a <b>termio</b> structure. Those terminal parameters that can be stored in a <b>termio</b> structure are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted. This form should be used when changing parameters that will affect output.
<b>TCSETAF</b>	The argument is a pointer to a <b>termio</b> structure. Those terminal parameters that can be stored in a <b>termio</b> structure are set from the values stored in that structure. The change occurs after all characters queued for output have been transmitted; all characters queued for input are discarded and then the change occurs.
<b>TCSBRK</b>	The argument is an <b>int</b> value. Wait for the output to drain. If the argument is 0, then send a break (zero-valued bits for 0.25 seconds). This define is available by <b>#include &lt;sys/termio.h&gt;</b>
<b>TCXONC</b>	Start/stop control. The argument is an <b>int</b> value. If the argument is <b>TCOOFF</b> (0), suspend output; if <b>TCOON</b> (1), restart suspended output; if <b>TCIOFF</b> (2), suspend input; if <b>TCION</b> (3), restart suspended input.
<b>TCFLSH</b>	The argument is an <b>int</b> value. If the argument is <b>TCIFLUSH</b> (0), flush the input queue; if <b>TCOFLUSH</b> (1), flush the output queue; if <b>TCIOFLUSH</b> (2), flush both the input and output queues.
<b>TIOCEXCL</b>	The argument is ignored. Exclusive-use mode is turned on; no further opens are permitted until the file has been closed, or a <b>TIOCNXCL</b> is issued. The default on open of a terminal file is that exclusive use mode is off. This <b>ioctl()</b> is only available by <b>#include &lt;sys/ttold.h&gt;</b> .
<b>TIOCNXCL</b>	The argument is ignored. Exclusive-use mode is turned off. This <b>ioctl()</b> is only available by <b>#include &lt;sys/ttold.h&gt;</b> .
<b>TIOCSCTTY</b>	The argument is an <b>int</b> . The system will attempt to assign the terminal as the caller's controlling terminal (see <b>The Controlling Terminal</b> above). If the caller is not the super-user and/or the argument is not 1, all of the normal permission checks apply. If the caller is the super-user and the argument is 1 the terminal will be assigned as the controlling terminal even if the terminal was currently in use as a controlling terminal by another session. <b>getty(8)</b> uses this method to acquire controlling terminals for <b>login(1)</b> because there exists a possibility that a daemon process may obtain the console before <b>getty(8)</b> .
<b>TIOCGPGRP</b>	The argument is a pointer to an <b>int</b> . Set the value of that <b>int</b> to the process group ID of the distinguished process group associated with the terminal. This call is allowed from a background process; however, the information may subsequently be changed by a foreground process. This <b>ioctl()</b> exists only for backward compatibility, use <b>tcgetpgrp(3V)</b> .
<b>TIOCSPGRP</b>	The argument is a pointer to an <b>int</b> . Associate the process group whose process group ID is specified by the value of that <b>int</b> with the terminal. The new process group value must be in the range of valid process group ID values, or it must be zero ("no process group"). Otherwise, the error <b>EINVAL</b> is returned. If any processes exist with a process ID or process group ID that is the same as the new process group value, then those processes must have the same real or saved user ID as the real or effective user ID of the calling process or be descendants of the calling process, or the effective user ID of the current process must be super-user. Otherwise, the error <b>EPERM</b> is returned. This <b>ioctl()</b> exists only for backward compatibility, use <b>tcsetpgrp()</b> , see <b>tcgetpgrp(3V)</b> .
<b>TIOCOUTQ</b>	The argument is a pointer to an <b>int</b> . Set the value of that <b>int</b> to the number of characters in the output stream that have not yet been sent to the terminal. This call is allowed from a background process.

- TIOCSTI** The argument is a pointer to a **char**. Pretend that character had been received as input.
- TIOCGWINSZ** The argument is a pointer to a **winsize** structure. The terminal driver's notion of the terminal size is stored into that structure. This call is allowed from a background process.
- TIOCSWINSZ** The argument is a pointer to a **winsize** structure. The terminal driver's notion of the terminal size is set from the values specified in that structure. If the new sizes are different from the old sizes, a **SIGWINCH** signal is sent to the process group of the terminal.
- TIOCMGET** The argument is a pointer to an **int**. The current state of the modem status lines is fetched and stored in the **int** pointed to by the argument. This call is allowed from a background process.
- TIOCMBIS** The argument is a pointer to an **int** whose value is a mask containing modem control lines to be turned on. The control lines whose bits are set in the argument are turned on; no other control lines are affected.
- TIOCMBIC** The argument is a pointer to an **int** whose value is a mask containing modem control lines to be turned off. The control lines whose bits are set in the argument are turned off; no other control lines are affected.
- TIOCMSET** The argument is a pointer to an **int** containing a new set of modem control lines. The modem control lines are turned on or off, depending on whether the bit for that mode is set or clear.
- TIOCGSOFTCAR** The argument is a pointer to an **int** whose value is 1 or 0, depending on whether the software carrier detect is turned on or off.
- TIOCSSOFTCAR** The argument is a pointer to an **int** whose value is 1 or 0. The value of the integer should be 0 to turn off software carrier, or 1 to turn it on.

**SEE ALSO**

**cs**(1), **login**(1), **stty**(1V), **fork**(2V), **getpgrp**(2V), **ioctl**(2), **open**(2V), **read**(2V), **sigvec**(2), **vfork**(2), **tcgetpgrp**(3V), **tty**(4), **ttytab**(5), **getty**(8), **init**(8), **ttysoftcar**(8)

**NAME**

tfs, TFS – translucent file service

**CONFIG**

*options*TFS

**SYNOPSIS**

```
#include <sys/mount.h>
mount("tfs", dir, M_NEWTYPE|flags, nfsargs);
```

**DESCRIPTION**

The translucent file service (TFS) supplies a copy-on-write filesystem allowing users to share file hierarchies while providing each user with a private hierarchy into which files are copied as they are modified. Consequently, users are isolated from each other's changes.

*nfsargs* specifies NFS style `mount(2V)` arguments, including the address of the file server (the `tfsd(8)`) and the file handle to be mounted. *dir* is the directory on which the TFS filesystem is to be mounted.

TFS allows a user to mount a private, writable filesystem in front of any number of public, read-only filesystems in such a way that the contents of the public filesystems remain visible behind the contents of the private filesystem. Any change made to a file that is being shared from a public filesystem will cause that file to be copied into the private filesystem, where the modification will be performed.

A directory in a TFS filesystem consists of a number of stacked directories. The searchpath TFS uses to look up a file in a directory corresponds to the stacking order: the TFS will search the "frontmost" directory first, then the directory behind it, and so on until the first occurrence of the file is found. Modifications to a file can be made only in the frontmost directory. TFS copies a file to the frontmost directory when the file is opened for writing with `open(2V)` or when its `stat(2V)` attributes are changed.

If a user removes a file which is not in the frontmost directory, TFS creates a *whiteout* entry in the frontmost directory and leaves the file intact in the back directory. This whiteout entry makes it appear that the file no longer exists, although the file can be reinstated in the directory by using the `unwhiteout(1)` command to remove the whiteout entry. The `lsw(1)` command lists whiteout entries.

TFS filesystems are served by the `tfsd(8)`. A TFS filesystem is mounted on a directory by making a `TFS_MOUNT` protocol request of the `tfsd`, specifying the directories that are to be stacked. The `tfsd` responds with a file handle, which the client then supplies to the `mount(2V)` system call, along with the address of the `tfsd`.

**SEE ALSO**

`lsw(1)`, `unwhiteout(1)`, `mount(2V)`, `tfsd(8)`, `mount_tfs(8)`

**NAME**

**timod** – Transport Interface cooperating STREAMS module

**CONFIG**

**pseudo-device tim64**

**DESCRIPTION**

**timod** is a STREAMS module for use with the Transport Interface (TI) functions of the Network Services library (see Section 3). The **timod** module converts a set of **ioctl(2)** calls into STREAMS messages that may be consumed by a transport protocol provider which supports the Transport Interface. This allows a user to initiate certain TI functions as atomic operations.

The **timod** module must be pushed onto only a *stream* terminated by a transport protocol provider which supports the TI.

All STREAMS messages, with the exception of the message types generated from the **ioctl()** commands described below, are transparently passed to the neighboring STREAMS module or driver. The messages generated from the following **ioctl()** commands are recognized and processed by the **timod** module. The format of the **ioctl()** call is:

Where, on issuance, **size** is the size of the appropriate TI message to be sent to the transport provider and on return **size** is the size of the appropriate TI message from the transport provider in response to the issued TI message. **buf** is a pointer to a buffer large enough to hold the contents of the appropriate TI messages. The TI message types are defined in `<sys/tihdr.h>`. The possible values for the **cmd** field are:

TI_BIND	Bind an address to the underlying transport protocol provider. The message issued to the <b>TI_BIND ioctl()</b> is equivalent to the TI message type <b>T_BIND_REQ</b> and the message returned by the successful completion of the <b>ioctl()</b> is equivalent to the TI message type <b>T_BIND_ACK</b> .
TI_UNBIND	Unbind an address from the underlying transport protocol provider. The message issued to the <b>TI_UNBIND ioctl()</b> is equivalent to the TI message type <b>T_UNBIND_REQ</b> and the message returned by the successful completion of the <b>ioctl()</b> is equivalent to the TI message type <b>T_OK_ACK</b> .
TI_GETINFO	Get the TI protocol specific information from the transport protocol provider. The message issued to the <b>TI_GETINFO ioctl()</b> is equivalent to the TI message type <b>T_INFO_REQ</b> and the message returned by the successful completion of the <b>ioctl()</b> is equivalent to the TI message type <b>T_INFO_ACK</b> .
TI_OPTMGMT	Get, set or negotiate protocol specific options with the transport protocol provider. The message issued to the <b>TI_OPTMGMT ioctl()</b> is equivalent to the TI message type <b>T_OPTMGMT_REQ</b> and the message returned by the successful completion of the <b>ioctl()</b> is equivalent to the TI message type <b>T_OPTMGMT_ACK</b> .

**SEE ALSO**

**tirdwr(4)**

*Network Programming*

**DIAGNOSTICS**

If the **ioctl()** system call returns with a value greater than 0, the lower 8 bits of the return value will be one of the TI error codes as defined in `<sys/tiuser.h>`. If the TI error is of type **TSYSERR**, then the next 8 bits of the return value will contain an error as defined in `<sys/errno.h>` (see **intro(2)**).

**NAME**

tirdwr – Transport Interface read/write interface STREAMS module

**CONFIG**

**pseudo-device tirw64**

**DESCRIPTION**

**tirdwr** is a STREAMS module that provides an alternate interface to a transport provider which supports the Transport Interface (TI) functions of the Network Services library (see Section 3). This alternate interface allows a user to communicate with the transport protocol provider using the **read(2V)** and **write(2V)** system calls. The **putmsg(2)** and **getmsg(2)** system calls may also be used. However, **putmsg()** and **getmsg()** can only transfer data messages between user and *stream*.

The **tirdwr** module must only be pushed (see **I\_PUSH** in **streamio(4)**) onto a **stream** terminated by a transport protocol provider which supports the TI. After the **tirdwr** module has been pushed onto a *stream*, none of the Transport Interface functions can be used. Subsequent calls to TI functions cause an error on the *stream*. Once the error is detected, subsequent system calls on the **stream** return an error with **errno** set to **EPROTO**.

The following are the actions taken by the **tirdwr** module when pushed on the **stream**, popped (see **I\_POP** in **streamio(4)**) off the *stream*, or when data passes through it.

**push** When the module is pushed onto a **stream**, it checks any existing data destined for the user to ensure that only regular data messages are present. It ignores any messages on the **stream** that relate to process management, such as messages that generate signals to the user processes associated with the *stream*. If any other messages are present, the **I\_PUSH** returns an error with **errno** set to **EPROTO**.

**write** The module takes the following actions on data that originated from a **write()** system call:

All messages with the exception of messages that contain control portions (see **putmsg(2)** and **getmsg(2)**) are transparently passed onto the module's downstream neighbor.

Any zero length data message is freed by the module and is not passed onto the module's downstream neighbor.

Any message with a control portion generates an error, and any further system calls associated with the **stream** fail with **errno** set to **EPROTO**.

**read** The module takes the following actions on data that originated from the transport protocol provider:

All messages with the exception of those that contain control portions (see the **putmsg** and **getmsg** system calls) are transparently passed onto the module's upstream neighbor.

The action taken on messages with control portions is as follows:

- Messages that represent expedited data generate an error. All further system calls associated with the **stream** fail with **errno** set to **EPROTO**.
- Any data messages with control portions have the control portions removed from the message prior to passing the message on to the upstream neighbor.
- Messages that represent an orderly release indication from the transport provider generate a zero length data message, indicating the end of file, which are sent to the reader of the *stream*. The orderly release message itself is freed by the module.
- Messages that represent an abortive disconnect indication from the transport provider cause all further **write()** and **putmsg()** calls to fail with **errno** set to **ENXIO**. All further **read()** and **getmsg()** calls return zero length data (indicating an EOF) once all previous data has been read.



- With the exception of the above rules, all other messages with control portions generate an error and all further system calls associated with the **stream** fail with **errno** set to **EPROTO**.

Any zero length data messages are freed by the module and they are not passed onto the module's upstream neighbor.

**pop** When the module is popped off the **stream** or the **stream** is closed, the module takes the following action:

If an orderly release indication has been previously received, then an orderly release request is sent to the remote side of the transport connection.

**SEE ALSO**

**intro(2)**, **getmsg(2)**, **putmsg(2)**, **read(2V)**, **write(2V)**, **intro(3)**, **streamio(4)**, **timod(4)**

*Network Programming*

**NAME**

**tm** – Tapemaster 1/2 inch tape controller

**CONFIG — SUN-3, SUN-3x SYSTEMS**

**controller tm0 at vme16d16 ? csr 0xa0 priority 3 vector tmintr 0x60**

**controller tm1 at vme16d16 ? csr 0xa2 priority 3 vector tmintr 0x61**

**tape mt0 at tm0 drive 0 flags 1**

**tape mt0 at tm1 drive 0 flags 1**

**DESCRIPTION**

The Tapemaster tape controller controls Pertec-interface 1/2" tape drives such as the CDC Keystone, providing a standard tape interface to the device, see **mtio(4)**. This controller supports single-density or speed drives.

The **tm** driver supports the character device interface. The driver returns an ENOTTY error on unsupported ioctls.

The **tm** driver does not support the backspace file to beginning of file (MTNBSF n) command. The equivalent positioning can be obtained by using MTBSF (n+1) followed by MTF SF 1.

Half-inch reel tape devices do not support the retension ioctl.

**FILES**

<b>/dev/rmt*</b>	rewinding
<b>/dev/nrmt*</b>	non-rewinding

**SEE ALSO**

**mt(1), tar(1), mtio(4), st(4S), xt(4S)**

**BUGS**

The Tapemaster controller does not provide for byte-swapping and the resultant system overhead prevents streaming transports from streaming.

The system should remember which controlling terminal has the tape drive open and write error messages to that terminal rather than on the console.

The Tapemaster controller is not supported on Sun-4 systems.

**WARNINGS**

The Tapemaster interface will not be supported in a future release. The Xylogics 472 controller and **xt** driver replace the Tapemaster controller and **tm** driver.

**NAME**

**tmpfs** – memory based filesystem

**CONFIG**

**options TMPFS**

**SYNOPSIS**

```
#include <sys/mount.h>
mount ("tmpfs", dir, M_NEWTYPE | flags, args);
```

**DESCRIPTION**

**tmpfs** is a memory based filesystem which uses kernel resources relating to the VM system and page cache as a filesystem. Once mounted, a **tmpfs** filesystem provides standard file operations and semantics. **tmpfs** is so named because files and directories are not preserved across reboot or unmounts, all files residing on a **tmpfs** filesystem that is unmounted will be lost.

**tmpfs** filesystems are mounted either with the command:

```
mount -t tmp swap directory-name
```

or by placing the line

```
swap directory-name tmp rw 0 0
```

in your **/etc/fstab** file and using the **mount(8)** command as normal. The **/etc/rc.local** file contains commands to mount a **tmpfs** filesystem on **/tmp** at multi-user startup time but is by default commented out. To mount a **tmpfs** filesystem on **/tmp** (maximizing possible performance improvements), add the above line to **/etc/fstab** and uncomment the following line in **/etc/rc.local**:

```
#mount /tmp
```

**tmpfs** is designed as a performance enhancement which is achieved by cacheing the writes to files residing on a **tmpfs** filesystem. Performance improvements are most noticeable when a large number of short lived files are written and accessed on a **tmpfs** filesystem. Large compilations with **tmpfs** mounted on **/tmp** are a good example of this.

Users of **tmpfs** should be aware of some tradeoffs involved in mounting a **tmpfs** filesystem. The resources used by **tmpfs** are the same as those used when commands are executed (for example, swap space allocation). This means that a large sized or number of **tmpfs** files can affect the amount of space left over for programs to execute. Likewise, programs requiring large amounts of memory use up the space available to **tmpfs**. Users running into these constraints (for example, running out of space on **tmpfs**) can allocate more swap space by using the **swapon(8)** command.

Normal filesystem writes are scheduled to be written to a permanent storage medium along with all control information associated with the file (for example, modification time, file permissions). **tmpfs** control information resides only in memory and never needs to be written to permanent storage. File data remains in core until memory demands are sufficient to cause pages associated with **tmpfs** to be reused at which time they are copied out to swap.

**SEE ALSO**

**df(1V)**, **mount(2V)**, **umount(2V)**, **fstab(5)**, **mount(8)**, **swapon(8)**

*System Services Overview,*  
*System and Network Administration*

**NOTES**

**swapon** to a **tmpfs** file is not supported.

**df(1V)** output is of limited accuracy since a **tmpfs** filesystem size is not static and the space available to **tmpfs** is dependent on the swap space demands of the entire system.

**DIAGNOSTICS**

If **tmpfs** runs out of space, one of the following messages will be printed to the console.

**directory: file system full, anon reservation exceeded**

**directory: file system full, anon allocation exceeded**

A page could not be allocated while writing to a file. This can occur if **tmpfs** is attempting to write more than it is allowed, or if currently executing programs are using a lot of memory. To make more space available, remove unnecessary files, exit from some programs, or allocate more swap space using **swapon(8)**.

**directory: file system full, kmem\_alloc failure**

**tmpfs** ran out of physical memory while attempting to create a new file or directory. Remove unnecessary files or directories or install more physical memory.

**WARNINGS**

A **tmpfs** filesystem should *not* be mounted on **/var/tmp**, this directory is used by **vi(1)** for preserved files. Files and directories on a **tmpfs** filesystem are not preserved across reboots or unmounts. Command scripts or programs which count on this will not work as expected.

**NAME**

`ttcompat` – V7 and 4BSD STREAMS compatibility module

**CONFIG**

None; included by default.

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/stream.h>
#include <sys/stropts.h>

ioctl(fd, I_PUSH, "ttcompat");
```

**DESCRIPTION**

`ttcompat` is a STREAMS module that translates the `ioctl` calls supported by the older Version 7 and 4BSD terminal drivers into the `ioctl` calls supported by the `termio(4)` interface. All other messages pass through this module unchanged; the behavior of `read` and `write` calls is unchanged, as is the behavior of `ioctl` calls other than the ones supported by `ttcompat`.

Normally, this module is automatically pushed onto a stream when a terminal device is opened; it does not have to be explicitly pushed onto a stream. This module requires that the `termio` interface be supported by the modules and driver downstream. The `TCGETS`, `TCSETS`, and `TCSETSF` `ioctl` calls must be supported; if any information set or fetched by those `ioctl` calls is not supported by the modules and driver downstream, some of the V7/4BSD functions may not be supported. For example, if the `CBAUD` bits in the `c_cflag` field are not supported, the functions provided by the `sg_ispeed` and `sg_ospeed` fields of the `sgttyb` structure (see below) will not be supported. If the `TIOCFUSH` `ioctl` is not supported, the function provided by the `TIOCFUSH` `ioctl` will not be supported. If the `TCXONC` `ioctl` is not supported, the functions provided by the `TIOCSTOP` and `TIOCSTART` `ioctl` calls will not be supported. If the `TIOCMBIS` and `TIOCMBCI` `ioctl` calls are not supported, the functions provided by the `TIOCSDTTR` and `TIOCCDTR` `ioctl` calls will not be supported.

The basic `ioctl` calls use the `sgttyb` structure defined by `<sys/ioctl.h>`:

```
struct sgttyb {
    char    sg_ispeed;
    char    sg_ospeed;
    char    sg_erase;
    char    sg_kill;
    short   sg_flags;
};
```

The `sg_ispeed` and `sg_ospeed` fields describe the input and output speeds of the device, and reflect the values in the `c_cflag` field of the `termio` structure. The `sg_erase` and `sg_kill` fields of the argument structure specify the erase and kill characters respectively, and reflect the values in the `VERASE` and `VKILL` members of the `c_cc` field of the `termio` structure.

The `sg_flags` field of the argument structure contains several flags that determine the system's treatment of the terminal. They are mapped into flags in fields of the terminal state, represented by the `termio` structure.

Delay type 0 is always mapped into the equivalent delay type 0 in the `c_oflag` field of the `termio` structure. Other delay mappings are performed as follows:

<code>sg_flags</code>	<code>c_oflag</code>
<code>BS1</code>	<code>BS1</code>
<code>FF1</code>	<code>VT1</code>
<code>CR1</code>	<code>CR2</code>
<code>CR2</code>	<code>CR3</code>
<code>CR3</code>	not supported
<code>TAB1</code>	<code>TAB1</code>
<code>TAB2</code>	<code>TAB2</code>

<b>XTABS</b>	<b>TAB3</b>
<b>NL1</b>	<b>ONLRET CR1</b>
<b>NL2</b>	<b>NL1</b>

If previous **TIOCLSET** or **TIOCLBIS ioctl** calls have not selected **LITOUT** or **PASS8** mode, and if **RAW** mode is not selected, the **ISTRIP** flag is set in the **c\_iflag** field of the **termio** structure, and the **EVENP** and **ODDP** flags control the parity of characters sent to the terminal and accepted from the terminal:

**0** Parity is not to be generated on output or checked on input; the character size is set to **CS8** and the **PARENB** flag is cleared in the **c\_cflag** field of the **termio** structure.

**EVENP** Even parity characters are to be generated on output and accepted on input; the **INPCK** flag is set in the **c\_iflag** field of the **termio** structure, the character size is set to **CS7** and the **PARENB** flag is set in the **c\_cflag** field of the **termio** structure.

**ODDP** Odd parity characters are to be generated on output and accepted on input; the **INPCK** flag is set in the **c\_iflag** field, the character size is set to **CS7** and the **PARENB** and **PARODD** flags are set in the **c\_cflag** field of the **termio** structure.

**EVENP|ODDP**

Even parity characters are to be generated on output and characters of either parity are to be accepted on input; the **INPCK** flag is cleared in the **c\_iflag** field, the character size is set to **CS7** and the **PARENB** flag is set in the **c\_cflag** field of the **termio** structure.

The **RAW** flag disables all output processing (the **OPOST** flag in the **c\_oflag** field, and the **XCASE** flag in the **c\_iflag** field, are cleared in the **termio** structure) and input processing (all flags in the **c\_iflag** field other than the **IXOFF** and **IXANY** flags are cleared in the **termio** structure). 8 bits of data, with no parity bit, are accepted on input and generated on output; the character size is set to **CS8** and the **PARENB** and **PARODD** flags are cleared in the **c\_cflag** field of the **termio** structure. The signal-generating and line-editing control characters are disabled by clearing the **ISIG** and **ICANON** flags in the **c\_lflag** field of the **termio** structure.

The **CRMOD** flag turn input **RETURN** characters into **NEWLINE** characters, and output and echoed **NEWLINE** characters to be output as a **RETURN** followed by a **LINEFEED**. The **ICRNL** flag in the **c\_iflag** field, and the **OPOST** and **ONLCR** flags in the **c\_oflag** field, are set in the **termio** structure.

The **LCASE** flag maps upper-case letters in the ASCII character set to their lower-case equivalents on input (the **IUCLC** flag is set in the **c\_iflag** field), and maps lower-case letters in the ASCII character set to their upper-case equivalents on output (the **OLCUC** flag is set in the **c\_oflag** field). Escape sequences are accepted on input, and generated on output, to handle certain ASCII characters not supported by older terminals (the **XCASE** flag is set in the **c\_lflag** field).

Other flags are directly mapped to flags in the **termio** structure:

<b>sg_flags</b>	flags in <b>termio</b> structure
<b>CBREAK</b>	complement of <b>ICANON</b> in <b>c_lflag</b> field
<b>ECHO</b>	<b>ECHO</b> in <b>c_lflag</b> field
<b>TANDEM</b>	<b>IXOFF</b> in <b>c_iflag</b> field

Another structure associated with each terminal specifies characters that are special in both the old Version 7 and the newer 4BSD terminal interfaces. The following structure is defined by `<sys/ioctl.h>`:

```

struct tchars {
    char    t_intrc;        /* interrupt */
    char    t_quitc;       /* quit */
    char    t_startc;      /* start output */
    char    t_stopc;       /* stop output */
    char    t_eofc;       /* end-of-file */
    char    t_brkc;       /* input delimiter (like nl) */
};

```

The characters are mapped to members of the `c_cc` field of the `termio` structure as follows:

tchars	c_cc index
<code>t_intrc</code>	<code>VINTR</code>
<code>t_quite</code>	<code>VQUIT</code>
<code>t_startc</code>	<code>VSTART</code>
<code>t_stopc</code>	<code>VSTOP</code>
<code>t_eofc</code>	<code>VEOF</code>
<code>t_brkc</code>	<code>VEOL</code>

Also associated with each terminal is a local flag word, specifying flags supported by the new 4BSD terminal interface. Most of these flags are directly mapped to flags in the `termio` structure:

local flags	flags in <code>termio</code> structure
<code>LCRTBS</code>	not supported
<code>LPRTERA</code>	<code>ECHOPRT</code> in the <code>c_lflag</code> field
<code>LCRTERA</code>	<code>ECHOE</code> in the <code>c_lflag</code> field
<code>LTILDE</code>	not supported
<code>LTOSTOP</code>	<code>TOSTOP</code> in the <code>c_lflag</code> field
<code>LFLUSHO</code>	<code>FLUSHO</code> in the <code>c_lflag</code> field
<code>LNOHANG</code>	<code>CLOCAL</code> in the <code>c_cflag</code> field
<code>LCRTKIL</code>	<code>ECHOKE</code> in the <code>c_lflag</code> field
<code>LCTLECH</code>	<code>CTLECH</code> in the <code>c_lflag</code> field
<code>LPENDIN</code>	<code>PENDIN</code> in the <code>c_lflag</code> field
<code>LDECCTQ</code>	complement of <code>IXANY</code> in the <code>c_iflag</code> field
<code>LNOFLSH</code>	<code>NOFLSH</code> in the <code>c_lflag</code> field

Another structure associated with each terminal is the `ltchars` structure which defines control characters for the new 4BSD terminal interface. Its structure is:

```

struct ltchars {
    char  t_suspc;      /* stop process signal */
    char  t_dsuspc;    /* delayed stop process signal */
    char  t_rprntc;    /* reprint line */
    char  t_flushc;    /* flush output (toggles) */
    char  t_werasc;    /* word erase */
    char  t_inxctc;    /* literal next character */
};

```

The characters are mapped to members of the `c_cc` field of the `termio` structure as follows:

ltchars	c_cc index
<code>t_suspc</code>	<code>VSUSP</code>
<code>t_dsuspc</code>	<code>VDSUSP</code>
<code>t_rprntc</code>	<code>VREPRINT</code>
<code>t_flushc</code>	<code>VDISCARD</code>
<code>t_werasc</code>	<code>VWERASE</code>
<code>t_inxctc</code>	<code>VLNEXT</code>

## IOCTLS

`ttcompat` responds to the following `ioctl` calls. All others are passed to the module below.

**TIOCGETP** The argument is a pointer to an `sgttyb` structure. The current terminal state is fetched; the appropriate characters in the terminal state are stored in that structure, as are the input and output speeds. The values of the flags in the `sg_flags` field are derived from the flags in the terminal state and stored in the structure.

- TIOCSETP** The argument is a pointer to an **sgttyb** structure. The appropriate characters and input and output speeds in the terminal state are set from the values in that structure, and the flags in the terminal state are set to match the values of the flags in the **sg\_flags** field of that structure. The state is changed with a **TCSETS*f* ioctl**, so that the interface delays until output is quiescent, then throws away any unread characters, before changing the modes.
- TIOCSETN** The argument is a pointer to an **sgttyb** structure. The terminal state is changed as **TIOCSETP** would change it, but a **TCSETS *ioctl*** is used, so that the interface neither delays nor discards input.
- TIOCHPCL** The argument is ignored. The **HUPCL** flag is set in the **c\_cflag** word of the terminal state.
- TIOCFLUSH** The argument is a pointer to an **int** variable. If its value is zero, all characters waiting in input or output queues are flushed. Otherwise, the value of the **int** is treated as the logical OR of the **FREAD** and **FWRITE** flags defined by **<sys/file.h>**; if the **FREAD** bit is set, all characters waiting in input queues are flushed, and if the **FWRITE** bit is set, all characters waiting in output queues are flushed.
- TIOCSBRK** The argument is ignored. The break bit is set for the device.
- TIOCCBRK** The argument is ignored. The break bit is cleared for the device.
- TIOCSDTR** The argument is ignored. The Data Terminal Ready bit is set for the device.
- TIOCCDTR** The argument is ignored. The Data Terminal Ready bit is cleared for the device.
- TIOCSTOP** The argument is ignored. Output is stopped as if the **STOP** character had been typed.
- TIOCSTART** The argument is ignored. Output is restarted as if the **START** character had been typed.
- TIOCGETC** The argument is a pointer to an **tchars** structure. The current terminal state is fetched, and the appropriate characters in the terminal state are stored in that structure.
- TIOCSETC** The argument is a pointer to an **tchars** structure. The values of the appropriate characters in the terminal state are set from the characters in that structure.
- TIOCLGET** The argument is a pointer to an **int**. The current terminal state is fetched, and the values of the local flags are derived from the flags in the terminal state and stored in the **int** pointed to by the argument.
- TIOCLBIS** The argument is a pointer to an **int** whose value is a mask containing flags to be set in the local flags word. The current terminal state is fetched, and the values of the local flags are derived from the flags in the terminal state; the specified flags are set, and the flags in the terminal state are set to match the new value of the local flags word.
- TIOCLBIC** The argument is a pointer to an **int** whose value is a mask containing flags to be cleared in the local flags word. The current terminal state is fetched, and the values of the local flags are derived from the flags in the terminal state; the specified flags are cleared, and the flags in the terminal state are set to match the new value of the local flags word.
- TIOCLSET** The argument is a pointer to an **int** containing a new set of local flags. The flags in the terminal state are set to match the new value of the local flags word.
- TIOCG LTC** The argument is a pointer to an **ltchars** structure. The values of the appropriate characters in the terminal state are stored in that structure.
- TIOCS LTC** The argument is a pointer to an **ltchars** structure. The values of the appropriate characters in the terminal state are set from the characters in that structure.

**SEE ALSO****ioctl(2), termio(4)**



**NAME**

tty – controlling terminal interface

**DESCRIPTION**

The file `/dev/tty` is, in each process, a synonym for the controlling terminal of that process, if any. It is useful for programs or shell sequences that wish to be sure of writing messages on the terminal no matter how output has been redirected. It can also be used for programs that demand the name of a file for output, when typed output is desired and it is tiresome to find out what terminal is currently in use.

**IOCTLS**

In addition to the `ioctl()` requests supported by the device that `tty` refers to, the following `ioctl()` request is supported:

**TIOCNOTTY**      Detach the current process from its controlling terminal, and remove it from its current process group, without attaching it to a new process group (that is, set its process group ID to zero). This `ioctl()` call only works on file descriptors connected to `/dev/tty`; this is used by daemon processes when they are invoked by a user at a terminal. The process attempts to open `/dev/tty`; if the open succeeds, it detaches itself from the terminal by using `TIOCNOTTY`, while if the open fails, it is obviously not attached to a terminal and does not need to detach itself.

**FILES**

`/dev/tty`

**SEE ALSO**

`termio(4)`

**NAME**

udp – Internet User Datagram Protocol

**SYNOPSIS**

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

s = socket(AF_INET, SOCK_DGRAM, 0);
```

**DESCRIPTION**

UDP is a simple, unreliable datagram protocol which is used to support the **SOCK\_DGRAM** abstraction for the Internet protocol family. It is layered directly above the Internet Protocol (IP). UDP sockets are connectionless, and are normally used with the **sendto**, **sendmsg**, **recvfrom**, and **recvmsg** system calls (see **send(2)** and **recv(2)**). If the **connect(2)** system call is used to fix the destination for future packets, then the **recv(2)** or **read(2V)** and **send(2)** or **write(2V)** system calls may be used.

UDP address formats are identical to those used by the Transmission Control Protocol (TCP). Like TCP, UDP uses a port number along with an IP address to identify the endpoint of communication. Note: the UDP port number space is separate from the TCP port number space (that is, a UDP port may not be “connected” to a TCP port). The **bind(2)** system call can be used to set the local address and port number of a UDP socket. The local IP address may be left unspecified in the **bind** call by using the special value **INADDR\_ANY**. If the **bind** call is not done, a local IP address and port number will be assigned to each packet as it is sent. Broadcast packets may be sent (assuming the underlying network supports this) by using a reserved “broadcast address”; this address is network interface dependent. Broadcasts may only be sent by the super-user.

Options at the IP level may be used with UDP; see **ip(4P)**.

There are a variety of ways that a UDP packet can be lost or discarded, including a failure of the underlying communication mechanism. UDP implements a checksum over the data portion of the packet. If the checksum of a received packet is in error, the packet will be dropped with no indication given to the user. A queue of received packets is provided for each UDP socket. This queue has a limited capacity. Arriving datagrams which will not fit within its *high-water* capacity are silently discarded.

UDP processes Internet Control Message Protocol (ICMP) error messages received in response to UDP packets it has sent. See **icmp(4P)**. ICMP “source quench” messages are ignored. ICMP “destination unreachable,” “time exceeded” and “parameter problem” messages disconnect the socket from its peer so that subsequent attempts to send packets using that socket will return an error. UDP will not guarantee that packets are delivered in the order they were sent. As well, duplicate packets may be generated in the communication process.

**ERRORS**

A socket operation may fail if:

<b>EISCONN</b>	A <b>connect</b> operation was attempted on a socket on which a <b>connect</b> operation had already been performed, and the socket could not be successfully disconnected before making the new connection.
<b>EISCONN</b>	A <b>sendto</b> or <b>sendmsg</b> operation specifying an address to which the message should be sent was attempted on a socket on which a <b>connect</b> operation had already been performed.
<b>ENOTCONN</b>	A <b>send</b> or <b>write</b> operation, or a <b>sendto</b> or <b>sendmsg</b> operation not specifying an address to which the message should be sent, was attempted on a socket on which a <b>connect</b> operation had not already been performed.
<b>EADDRINUSE</b>	A <b>bind</b> operation was attempted on a socket with a network address/port pair that has already been bound to another socket.
<b>EADDRNOTAVAIL</b>	A <b>bind</b> operation was attempted on a socket with a network address for which no network interface exists.

**EINVAL** A **sendmsg** operation with a non-NULL **msg\_accrights** was attempted.

**EACCES** A **bind** operation was attempted with a "reserved" port number and the effective user ID of the process was not super-user.

**ENOBUFS** The system ran out of memory for internal data structures.

**SEE ALSO**

**bind(2), connect(2), read(2V), recv(2), send(2), write(2V), icmp(4P), inet(4F), ip(4P), tcp(4P)**

Postel, Jon, *User Datagram Protocol*, RFC 768, Network Information Center, SRI International, Menlo Park, Calif., August 1980. (Sun 800-1054-01)

**BUGS**

**SIOCSHIWAT** and **SIOCGHIWAT** **ioctl**'s to set and get the high water mark for the socket queue, and so that it can be changed from 2048 bytes to be larger or smaller, have been defined (in **sys/ioctl.h**) but not implemented.

Something sensible should be done with ICMP source quench error messages if the socket is bound to a peer socket.

**NAME**

unix – UNIX domain protocol family

**DESCRIPTION**

The Unix Domain protocol family provides support for socket-based communication between processes running on the local host. While both **SOCK\_STREAM** and **SOCK\_DGRAM** types are supported, the **SOCK\_STREAM** type often provides faster performance. Pipes, for instance, are built on Unix Domain **SOCK\_STREAM** sockets.

Unix Domain **SOCK\_DGRAM** sockets (also called datagram sockets) exist primarily for reasons of orthogonality under the BSD socket model. However, the overhead of reading or writing data is higher for the (connectionless) datagram sockets.

Unix Domain addresses are pathnames. In other words, two independent processes can communicate by specifying the same pathname as their communications rendezvous point. The **bind(2)** operation creates a special entry in the file system of type socket. If that pathname already exists (as a socket from a previous **bind()** operation, or as some other file system type), **bind()** will fail.

Sockets in the Unix domain protocol family use the following addressing structure:

```

struct sockaddr_un {
    short  sun_family;
    u_short sun_path[108];
};

```

To create or reference a Unix Domain socket, the **sun\_family** field should be set to **AF\_UNIX** and the **sun\_path** array should contain the path name of a rendezvous point.

Although Unix Domain sockets are faster than Internet Domain sockets for communication between local processes, the advantage of the additional flexibility afforded by the latter may outweigh performance issues. Where inter-process communication throughput is critical, a shared memory approach may be preferred.

Since there are no protocol families associated with Unix Domain sockets, the protocol argument to **socket(2)** should be zero.

When setting up a Unix Domain socket, the *length* argument to the **bind()** call is the amount of space within the **sockaddr\_un** structure, not including the pathname delimiter. One way to specify the length is:

```

sizeof(addr.sun_family) + strlen(path)

```

where *addr* is a structure of type **sockaddr\_un**, and *path* is a pointer to the pathname.

The limit of 108 characters is an artifact of the implementation.

Since closing a Unix Domain socket does not make the file system entry go away, an application should remove the entry using **unlink(2V)**, when finished.

**SEE ALSO**

**bind(2)**, **socket(2)**, **unlink(2V)**

*Network Programming*

**NAME**

vd – loadable modules interface

**CONFIG**

None; included with **options VDDRV**

**DESCRIPTION**

This pseudo-device provides kernel support for loadable modules. It is used exclusively by the **modload(8)**, **modunload(8)**, and **modstat(8)** utilities. Other programs should not use it.

**FILES**

**/dev/vd**

**SEE ALSO**

**modload()**, **modunload()**, **modstat()**

**WARNINGS**

The interface provided by **vd** is subject to change without notice.

**NAME**

vpc – Systech VPC-2200 Versatec printer/plotter and Centronics printer interface

**CONFIG**

```
device vpc0 at vme16d16 ? csr 0x480 priority 2 vector vpcintr 0x80
device vpc1 at vme16d16 ? csr 0x500 priority 2 vector vpcintr 0x81
```

**AVAILABILITY**

Sun-3, Sun-3/80 and Sun-4 systems only.

**DESCRIPTION**

This Sun interface to the Versatec printer/plotter and to Centronics printers is supported by the Systech parallel interface board, an output-only byte-wide DMA device. The device has one channel for Versatec devices and one channel for Centronics devices, with an optional long lines interface for Versatec devices.

Devices attached to this interface are normally handled by the line printer spooling system and should not be accessed by the user directly.

Opening the device `/dev/vpc0` or `/dev/lp0` may yield one of two errors: ENXIO indicates that the device is already in use; EIO indicates that the device is offline.

The Versatec printer/plotter operates in either print or plot mode. To set the printer into plot mode you should include `<sys/vcmd.h>` and use the `ioctl(2)` call:

```
ioctl(f, VSETSTATE, plotmd);
```

where `plotmd` is defined to be

```
int plotmd[] = { VPLOT, 0, 0 };
```

When going back into print mode from plot mode you normally eject paper by sending it an EOT after putting into print mode:

```
int prtmd[] = { VPRINT, 0, 0 };
```

```
...
```

```
fflush(vpc);
```

```
f = fileno(vpc);
```

```
ioctl(f, VSETSTATE, prtmd);
```

```
write(f, "\04", 1);
```

**FILES**

`/dev/vpc0`

`/dev/lp0`

**SEE ALSO**

`ioctl(2)`, `setbuf(3V)`

**BUGS**

If you use the standard I/O library on the Versatec, be sure to explicitly set a buffer using `setbuf(3V)`, since the library will not use buffered output by default, and will run very slowly.

**NAME**

win – Sun window system

**CONFIG**

**pseudo-device** *winnumber*  
**pseudo-device** *dtopnumber*

**DESCRIPTION**

The **win** pseudo-device accesses the system drivers supporting the Sun window system. *number*, in the device description line above, indicates the maximum number of windows supported by the system. *number* is set to 128 in the GENERIC system configuration file used to generate the kernel used in Sun systems as they are shipped. The *dtop* pseudo-device line indicates the number of separate “desktops” (frame buffers) that can be actively running the Sun window system at once. In the GENERIC file, this number is set to 4.

Each window in the system is represented by a **/dev/win\*** device. The windows are organized as a tree with windows being subwindows of their parents, and covering/covered by their siblings. Each window has a position in the tree, a position on a display screen, an input queue, and information telling what parts of it are exposed.

The window driver multiplexes keyboard and mouse input among the several windows, tracks the mouse with a cursor on the screen, provides each window access to information about what parts of it are exposed, and notifies the manager process for a window when the exposed area of the window changes so that the window may repair its display.

Full information on the window system functions is given in the *SunView System Programmer's Guide*.

**FILES**

**/dev/win[0-9]**  
**/dev/win[0-9][0-9]**

**SEE ALSO**

*SunView System Programmer's Guide*

**NAME**

xd – Disk driver for Xylogics 7053 SMD Disk Controller

**CONFIG — SUN-3, SUN-3x, SUN-4 SYSTEMS**

```

controller xdc0 at vme16d32 ? csr 0xee80 priority 2 vector xdintr 0x44
controller xdc1 at vme16d32 ? csr 0xee90 priority 2 vector xdintr 0x45
controller xdc2 at vme16d32 ? csr 0xeea0 priority 2 vector xdintr 0x46
controller xdc3 at vme16d32 ? csr 0xeeb0 priority 2 vector xdintr 0x47
disk xd0 at xdc0 drive 0
disk xd1 at xdc0 drive 1
disk xd2 at xdc0 drive 2
disk xd3 at xdc0 drive 3
disk xd4 at xdc1 drive 0
disk xd5 at xdc1 drive 1
disk xd6 at xdc1 drive 2
disk xd7 at xdc1 drive 3
disk xd8 at xdc2 drive 0
disk xd9 at xdc2 drive 1
disk xd10 at xdc2 drive 2
disk xd11 at xdc2 drive 3
disk xd12 at xdc3 drive 0
disk xd13 at xdc3 drive 1
disk xd14 at xdc3 drive 2
disk xd15 at xdc3 drive 3

```

The four **controller** lines given in the synopsis section above specify the first, second, third, and fourth Xylogics 7053 SMD disk controller in a Sun system.

**DESCRIPTION**

Files with minor device numbers 0 through 7 refer to various portions of drive 0; minor devices 8 through 15 refer to drive 1, and so on. The standard device names begin with **xd** followed by the drive number and then a letter a-h for partitions 0-7 respectively. The character ? stands here for a drive number in the range 0-7.

The block files access the disk using the system's normal buffering mechanism and may be read and written without regard to physical disk records. There is also a "raw" interface which provides for direct transmission between the disk and the user's read or write buffer. A single read or write call usually results in only one I/O operation; therefore raw I/O is considerably more efficient when many words are transmitted. The names of the raw files conventionally begin with an extra **r**.

In raw I/O counts should be a multiple of 512 bytes (a disk sector). Likewise **directory(3V)** calls should specify a multiple of 512 bytes.

If **flags 0x1** is specified, the overlapped seeks feature for that drive is turned off. Note: to be effective, the flag must be set on all drives for a specific controller. This action is necessary for controllers with older firmware, which have bugs preventing overlapped seeks from working properly.

**DISK SUPPORT**

This driver handles all SMD drives by reading a label from sector 0 of the drive which describes the disk geometry and partitioning.

The **xd?a** partition is normally used for the root file system on a disk, the **xd?b** partition as a paging area, and the **xd?c** partition for pack-pack copying (it normally maps the entire disk). The rest of the disk is normally the **xd?g** partition.

**FILES**

```

/dev/xd[0-7][a-h]    block files
/dev/rxd[0-7][a-h]  raw files

```



**SEE ALSO**

**lseek(2V), read(2V), write(2V), directory(3V), dkio(4S)**

**DIAGNOSTICS****xdcn: self test error**

Self test error in controller, see the Maintenance and Reference Manual.

**xdn: unable to read bad sector**

The bad sector forwarding information for the disk could not be read.

**xdn: initialization failed**

The drive could not be successfully initialized.

**xdn: unable to read label**

The drive geometry/partition table information could not be read.

**xdn: Corrupt label**

The geometry/partition label checksum was incorrect.

**xdn: offline**

A drive ready status is no longer detected, so the unit has been logically removed from the system. If the drive ready status is restored, the unit will automatically come back online the next time it is accessed.

**xdnc: cmd how (msg) blk #n abs blk #n**

A command such as read or write encountered an error condition (*how*): either it *failed*, the controller was *reset*, the unit was *restored*, or an operation was *retry*'ed. The *msg* is derived from the error number given by the controller, indicating a condition such as "drive not ready(rq, "sector not found" or "disk write protected". The *blk #* is the sector in error relative to the beginning of the partition involved. The *abs blk #* is the absolute block number of the sector in error. Some fields of the error message may be missing since the information is not always available.

**BUGS**

In raw I/O **read(2V)** and **write(2V)** truncate file offsets to 512-byte block boundaries, and **write(2V)** scribbles on the tail of incomplete blocks. Thus, in programs that are likely to access raw devices, **read(2V)**, **write(2V)** and **lseek(2V)** should always deal in 512-byte multiples.

Older revisions of the firmware do not properly support overlapped seeks. This will only affect systems with multiple disks on a single controller. If a large number of "zero sector count" errors appear, you should use the **flags** field to disable overlapped seeks.

**NAME**

xt – Xylogics 472 1/2 inch tape controller

**CONFIG — SUN-3, SUN-4 SYSTEMS**

controller xtc0 at vme16d16 ? csr 0xee60 priority 3 vector xtintr 0x64

controller xtc1 at vme16d16 ? csr 0xee68 priority 3 vector xtintr 0x65

tape xt0 at xtc0 drive 0 flags 1

tape xt1 at xtc1 drive 0 flags 1

**DESCRIPTION**

The Xylogics 472 tape controller controls Pertec-interface 1/2" tape drives such as the Fujitsu M2444 and the CDC Keystone III, providing a standard tape interface to the device see **mtio(4)**. This controller is used to support high speed or high density drives, which are not supported effectively by the older Tapemaster controller (see **tm(4S)**).

The flags field is used to control remote density select operation: a 0 specifies no remote density selection is to be attempted, a 1 specifies that the Pertec density-select line is used to toggle between high and low density; a 2 specifies that the Pertec speed-select line is used to toggle between high and low density. The default is 1, which is appropriate for the Fujitsu M2444, the CDC Keystone III (92185) and the Telex 9250. In no case will the controller select among more than 2 densities.

The xt driver supports the character device interface.

**EOT Handling**

The user will be notified of end of tape (EOT) on write by a 0 byte count returned the first time this is attempted. This write must be retried by the user. Subsequent writes will be successful until the tape winds off the reel. Read past EOT is transparent to the user.

**Ioctls**

Not all devices support all ioctls. The driver returns an ENOTTY error on unsupported ioctls.

1/2" tape devices do not support the tape retension function.

**FILES**

<b>/dev/rmt0</b>	low density operation, typically 1600 bpi
<b>/dev/rmt8</b>	high density operation, typically 6250 bpi
<b>/dev/nrmt*</b>	non-rewinding

**SEE ALSO**

**mt(1)**, **tar(1)**, **mtio(4)**, **st(4S)**, **suninstall(8)**

**BUGS**

Record sizes are restricted to an even number of bytes.

Absolute file positioning is not fully supported; it is only meant to be used by **suninstall(8)**.

**NAME**

xy – Disk driver for Xylogics 450 and 451 SMD Disk Controllers

**CONFIG — SUN-3, SUN-3x, SUN-4 SYSTEMS**

**controller xyc0 at vme16d16 ? csr 0xee40 priority 2 vector xyintr 0x48**

**controller xyc1 at vme16d16 ? csr 0xee48 priority 2 vector xyintr 0x49**

**disk xy0 at xyc0 drive 0**

**disk xy1 at xyc0 drive 1**

**disk xy2 at xyc1 drive 0**

**disk xy3 at xyc1 drive 1**

The two **controller** lines given in the synopsis sections above specify the first and second Xylogics 450 or 451 SMD disk controller in a Sun system.

**DESCRIPTION**

Files with minor device numbers 0 through 7 refer to various portions of drive 0; minor devices 8 through 15 refer to drive 1, and so on. The standard device names begin with xy followed by the drive number and then a letter a-h for partitions 0-7 respectively. The character '?' stands here for a drive number in the range 0-7.

The block files access the disk using the system's normal buffering mechanism and may be read and written without regard to physical disk records. There is also a "raw" interface which provides for direct transmission between the disk and the user's read or write buffer. A single read or write call usually results in only one I/O operation; therefore raw I/O is considerably more efficient when many words are transmitted. The names of the raw files conventionally begin with an extra r.

When using raw I/O, transfer counts should be multiples of 512 bytes (the size of a disk sector). Likewise, when using **lseek(2V)** to specify block offsets from which to perform raw I/O, the logical offset should also be a multiple of 512 bytes.

Due to word ordering differences between the disk controller and Sun computers, user buffers that are used for raw I/O must not begin on odd byte boundaries.

If **flags 0x1** is specified, the overlapped seeks feature for that drive is turned off. Note: to be effective, the flag must be set on all drives for a specific controller. This action is necessary for controllers with older firmware, which have bugs preventing overlapped seeks from working properly.

**DISK SUPPORT**

This driver handles all SMD drives by reading a label from sector 0 of the drive which describes the disk geometry and partitioning.

The **xy?a** partition is normally used for the root file system on a disk, the **xy?b** partition as a paging area, and the **xy?c** partition for pack-pack copying (it normally maps the entire disk). The rest of the disk is normally the **xy?g** partition.

**FILES**

<b>/dev/xy[0-7][a-h]</b>	block files
<b>/dev/rxy[0-7][a-h]</b>	raw files

**SEE ALSO**

**lseek(2V), read(2V), directory(3V), write(2V), dkio(4S)**

**DIAGNOSTICS**

**xycn : self test error**

Self test error in controller, see the Maintenance and Reference Manual.

**xycn: WARNING: n bit addresses**

The controller is strapped incorrectly. Sun systems use 20-bit addresses for Multibus based systems and 24-bit addresses for VMEbus based systems.

**xyn : unable to read bad sector info**

The bad sector forwarding information for the disk could not be read.

**xyn and xyn are of same type (n) with different geometries.**

The 450 and 451 do not support mixing the drive types found on these units on a single controller.

**xyn : initialization failed**

The drive could not be successfully initialized.

**xyn : unable to read label**

The drive geometry/partition table information could not be read.

**xyn : Corrupt label**

The geometry/partition label checksum was incorrect.

**xyn : offline**

A drive ready status is no longer detected, so the unit has been logically removed from the system. If the drive ready status is restored, the unit will automatically come back online the next time it is accessed.

**xync: cmd how (msg) blk #n abs blk #n**

A command such as read or write encountered an error condition (*how*): either it *failed*, the controller was *reset*, the unit was *restored*, or an operation was *retry*'ed. The *msg* is derived from the error number given by the controller, indicating a condition such as "drive not ready", "sector not found" or "disk write protected". The *blk #* is the sector in error relative to the beginning of the partition involved. The *abs blk #* is the absolute block number of the sector in error. Some fields of the error message may be missing since the information is not always available.

## BUGS

In raw I/O **read(2V)** and **write(2V)** truncate file offsets to 512-byte block boundaries, and **write(2V)** scribbles on the tail of incomplete blocks. Thus, in programs that are likely to access raw devices, **read(2V)**, **write(2V)** and **lseek(2V)** should always deal in 512-byte multiples.

Older revisions of the firmware do not properly support overlapped seeks. This will only affect systems with multiple disks on a single controller. If a large number of "zero sector count" errors appear, you should use the **flags** field to disable overlapped seeks.

**NAME**

zero – source of zeroes

**SYNOPSIS**

None; included with standard system.

**DESCRIPTION**

A zero special file is a source of zeroed unnamed memory.

Reads from a zero special file always return a buffer full of zeroes. The file is of infinite length.

Writes to a zero special file are always successful, but the data written is ignored.

Mapping a zero special file creates a zero-initialized unnamed memory object of a length equal to the length of the mapping and rounded up to the nearest page size as returned by `getpagesize(2)`. Multiple processes can share such a zero special file object provided a common ancestor mapped the object `MAP_SHARED`.

**FILES**

`/dev/zero`

**SEE ALSO**

`fork(2V)`, `getpagesize(2)`, `mmap(2)`

**NAME**

zs – Zilog 8530 SCC serial communications driver

**CONFIG — SUN-3 SYSTEM**

device zs0 at obio ? csr 0x20000 flags 3 priority 3  
device zs1 at obio ? csr 0x00000 flags 0x103 priority 3

**CONFIG — SUN-3x SYSTEM**

device zs0 at obio ? csr 0x62002000 flags 3 priority 3  
device zs1 at obio ? csr 0x62000000 flags 0x103 priority 3

**CONFIG — SUN-4 SYSTEM**

device zs1 at obio ? csr 0xf0000000 flags 0x103 priority 3  
device zs2 at obio 3 csr 0xe0000000 flags 3 priority 3

**CONFIG — SPARCSTATION 1 SYSTEM**

device-driver zs

**CONFIG — Sun386i SYSTEM**

device zs0 at obmem ? csr 0xFC000000 flags 3 irq 9 priority 6  
device zs1 at obmem ? csr 0xA0000020 flags 0x103 irq 9 priority 6

**SYNOPSIS**

```
#include <fcntl.h>
#include <sys/termios.h>
open("/dev/tty $n$ ", mode);
open("/dev/ttyd $n$ ", mode);
open("/dev/cuan", mode);
```

**DESCRIPTION**

The Zilog 8530 provides 2 serial communication ports with full modem control in asynchronous mode. Each port supports those `termio(4)` device control functions specified by flags in the `c_cflag` word of the `termios` structure and by the `IGNBRK`, `IGNPAR`, `PARMRK`, or `INPCK` flags in the `c_iflag` word of the `termios` structure are performed by the `zs` driver. All other `termio(4)` functions must be performed by STREAMS modules pushed atop the driver; when a device is opened, the `ldterm(4M)` and `ttcompat(4M)` STREAMS modules are automatically pushed on top of the stream, providing the standard `termio(4)` interface.

Of the synopsis lines above, the line for `zs0` specifies the serial I/O port(s) provided by the CPU board, the line for `zs1` specifies the Video Board ports (which are used for keyboard and mouse), the lines for `zs2` and `zs3` specify the first and second ports on the first SCSI board in a system, and those for `zs4` and `zs5` specify the first and second ports provided by the second SCSI board in a system, respectively.

Bit  $i$  of `flags` may be specified to say that a line is not properly connected, and that the line  $i$  should be treated as hard-wired with carrier always present. Thus specifying `flags 0x2` in the specification of `zs0` would treat line `/dev/ttyb` in this way.

Minor device numbers in the range 0 – 11 correspond directly to the normal tty lines and are named `/dev/ttya` and `/dev/ttyb` for the two serial ports on the CPU board and `/dev/ttys $n$`  for the ports on the SCSI boards;  $n$  is 0 or 1 for the ports on the first SCSI board, and 2 or 3 for the ports on the second SCSI board.

To allow a single tty line to be connected to a modem and used for both incoming and outgoing calls, a special feature, controlled by the minor device number, has been added. Minor device numbers in the range 128 – 139 correspond to the same physical lines as those above (that is, the same line as the minor device number minus 128).

A dial-in line has a minor device in the range 0 – 11 and is conventionally renamed `/dev/ttyd $n$` , where  $n$  is a number indicating which dial-in line it is (so that `/dev/ttyd0` is the first dial-in line), and the dial-out line corresponding to that dial-in line has a minor device number 128 greater than the minor device number of the dial-in line and is conventionally named `/dev/cuan`, where  $n$  is the number of the dial-in line.

The `/dev/cuan` lines are special in that they can be opened even when there is no carrier on the line. Once a `/dev/cuan` line is opened, the corresponding tty line can not be opened until the `/dev/cuan` line is closed; a blocking open will wait until the `/dev/cuan` line is closed (which will drop Data Terminal Ready, after which Carrier Detect will usually drop as well) and carrier is detected again, and a non-blocking open will return an error. Also, if the `/dev/ttydn` line has been opened successfully (usually only when carrier is recognized on the modem) the corresponding `/dev/cuan` line can not be opened. This allows a modem to be attached to e.g. `/dev/ttyd0` (renamed from `/dev/ttya`) and used for dial-in (by enabling the line for login in `/etc/ttytab`) and also used for dial-out (by `tip(1C)` or `uucp(1C)`) as `/dev/cua0` when no one is logged in on the line. Note: the bit in the `flags` word in the configuration file (see above) must be zero for this line, which enables hardware carrier detection.

#### IOCTLS

The standard set of `termio ioctl()` calls are supported by `zs`.

If the `CRTSCTS` flag in the `c_cflag` is set, output will be generated only if CTS is high; if CTS is low, output will be frozen. If the `CRTSCTS` flag is clear, the state of CTS has no effect. Breaks can be generated by the `TCSBRK`, `TIOCSBRK`, and `TIOCCBRK ioctl()` calls. The modem control lines `TIOCM_CAR`, `TIOCM_CTS`, `TIOCM_RTS`, and `TIOCM_DTR` are provided.

The input and output line speeds may be set to any of the speeds supported by `termio`. The speeds cannot be set independently; when the output speed is set, the input speed is set to the same speed.

#### ERRORS

An `open()` will fail if:

<code>ENXIO</code>	The unit being opened does not exist.
<code>EBUSY</code>	The dial-out device is being opened and the dial-in device is already open, or the dial-in device is being opened with a no-delay open and the dial-out device is already open.
<code>EBUSY</code>	The unit has been marked as exclusive-use by another process with a <code>TIOCEXCL ioctl()</code> call.
<code>EINTR</code>	The open was interrupted by the delivery of a signal.

#### FILES

<code>/dev/tty{a,b,s[0-3]}</code>	hardwired tty lines
<code>/dev/ttyd[0-9a-f]</code>	dial-in tty lines
<code>/dev/cua[0-9a-f]</code>	dial-out tty lines

#### SEE ALSO

`tip(1C)`, `uucp(1C)`, `mcp(4S)`, `mti(4S)`, `termio(4)`, `ldterm(4M)`, `ttcompat(4M)`, `ttysoftcar(8)`

#### DIAGNOSTICS

**zsn c: silo overflow.**

The 8530 character input silo overflowed before it could be serviced.

**zsn c: ring buffer overflow.**

The driver's character input ring buffer overflowed before it could be serviced.

---

Notes