



man Pages(2): System Calls

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900
U.S.A.

Part No: 805-3174-10
October 1998

Copyright 1998 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, SunDocs, Java, the Java Coffee Cup logo, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 1998 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, Californie 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, SunDocs, Java, le logo Java Coffee Cup, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Contents

PREFACE ix

Intro(2) 2

access(2) 36

acct(2) 38

acl(2) 40

adjtime(2) 42

alarm(2) 44

audit(2) 46

auditon(2) 47

auditsvc(2) 52

brk(2) 54

chdir(2) 56

chmod(2) 58

chown(2) 62

chroot(2) 65

close(2) 67

creat(2) 69

dup(2) 72

exec(2) 73

execl(2) 73
exit(2) 78
fcntl(2) 81
fork(2) 90
fpathconf(2) 94
getaudit(2) 97
getaudit(2) 99
getcontext(2) 101
getdents(2) 103
getgroups(2) 105
getitimer(2) 106
getmsg(2) 110
getpid(2) 113
getrlimit(2) 114
getsid(2) 118
getuid(2) 119
ioctl(2) 120
kill(2) 122
link(2) 124
llseek(2) 126
lseek(2) 127
_lwp_cond_signal(2) 129
_lwp_cond_wait(2) 130
_lwp_create(2) 132
_lwp_exit(2) 134
_lwp_info(2) 135
_lwp_kill(2) 136
_lwp_makecontext(2) 137

_lwp_mutex_lock(2) 138
_lwp_self(2) 139
_lwp_sema_wait(2) 140
_lwp_setprivate(2) 142
_lwp_suspend(2) 143
_lwp_wait(2) 144
memcntl(2) 145
mincore(2) 149
mkdir(2) 150
mknod(2) 152
mmap(2) 155
mount(2) 161
mprotect(2) 163
msgctl(2) 165
msgget(2) 167
msgrcv(2) 169
msgsnd(2) 171
munmap(2) 173
nice(2) 174
ntp_adjtime(2) 175
ntp_gettime(2) 176
open(2) 177
pause(2) 186
pcsample(2) 187
pipe(2) 189
poll(2) 190
p_online(2) 193
prctl(2) 196

prctlset(2) 207
processor_bind(2) 209
processor_info(2) 211
profil(2) 212
pset_bind(2) 214
pset_create(2) 216
pset_info(2) 218
ptrace(2) 219
putmsg(2) 221
read(2) 224
readlink(2) 230
rename(2) 232
resolvepath(2) 235
rmdir(2) 237
semctl(2) 239
semget(2) 242
semop(2) 244
setpgid(2) 247
setpgrp(2) 249
setregid(2) 250
setreuid(2) 251
setsid(2) 252
setuid(2) 253
shmctl(2) 255
shmget(2) 257
shmop(2) 259
sigaction(2) 262
sigaltstack(2) 265

_signotifywait(2) 267
sigpending(2) 269
sigprocmask(2) 270
sigsend(2) 272
sigsuspend(2) 274
sigwait(2) 276
__sparc_utrap_install(2) 280
stat(2) 286
statvfs(2) 290
stime(2) 293
swapctl(2) 294
symlink(2) 298
sync(2) 300
sysfs(2) 301
sysinfo(2) 302
time(2) 306
times(2) 307
uadmin(2) 309
ulimit(2) 311
umask(2) 313
umount(2) 314
uname(2) 315
unlink(2) 316
ustat(2) 318
utime(2) 319
utimes(2) 321
vfork(2) 323
vhangup(2) 325

wait(2) 326
waitid(2) 328
waitpid(2) 330
write(2) 332
yield(2) 339
Index 340

PREFACE

Overview

A man page is provided for both the naive user, and sophisticated user who is familiar with the SunOS operating system and is in need of on-line information. A man page is intended to answer concisely the question “What does it do?” The man pages in general comprise a reference manual. They are not intended to be a tutorial.

The following contains a brief description of each section in the man pages and the information it references:

- Section 1 describes, in alphabetical order, commands available with the operating system.
- Section 1M describes, in alphabetical order, commands that are used chiefly for system maintenance and administration purposes.
- Section 2 describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible returned value.
- Section 3 describes functions found in various libraries, other than those functions that directly invoke UNIX system primitives, which are described in Section 2 of this volume.
- Section 4 outlines the formats of various files. The C structure declarations for the file formats are given where applicable.
- Section 5 contains miscellaneous documentation such as character set tables.
- Section 6 contains available games and demos.

- Section 7 describes various special files that refer to specific hardware peripherals, and device drivers. STREAMS software drivers, modules and the STREAMS-generic set of system calls are also described.
- Section 9 provides reference information needed to write device drivers in the kernel operating systems environment. It describes two device driver interface specifications: the Device Driver Interface (DDI) and the Driver/Kernel Interface (DKI).
- Section 9E describes the DDI/DKI, DDI-only, and DKI-only entry-point routines a developer may include in a device driver.
- Section 9F describes the kernel functions available for use by device drivers.
- Section 9S describes the data structures used by drivers to share information between the driver and the kernel.

Below is a generic format for man pages. The man pages of each manual section generally follow this order, but include only needed headings. For example, if there are no bugs to report, there is no BUGS section. See the `intro` pages for more information and detail about each section, and `man(1)` for more information about man pages in general.

NAME	This section gives the names of the commands or functions documented, followed by a brief description of what they do.
SYNOPSIS	<p>This section shows the syntax of commands or functions. When a command or file does not exist in the standard path, its full pathname is shown. Options and arguments are alphabetized, with single letter arguments first, and options with arguments next, unless a different argument order is required.</p> <p>The following special characters are used in this section:</p> <p>[] The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.</p> <p>. . . Ellipses. Several values may be provided for the previous argument, or the previous argument can be specified multiple times, for example, ‘<code>filename . . .</code>’.</p>

| Separator. Only one of the arguments separated by this character can be specified at time.

{ } Braces. The options and/or arguments enclosed within braces are interdependent, such that everything enclosed must be treated as a unit.

PROTOCOL

This section occurs only in subsection 3R to indicate the protocol description file.

DESCRIPTION

This section defines the functionality and behavior of the service. Thus it describes concisely what the command does. It does not discuss OPTIONS or cite EXAMPLES.. Interactive commands, subcommands, requests, macros, functions and such, are described under USAGE.

IOCTL

This section appears on pages in Section 7 only. Only the device class which supplies appropriate parameters to the ioctl (2) system call is called `ioctl` and generates its own heading. `ioctl` calls for a specific device are listed alphabetically (on the man page for that specific device). `ioctl` calls are used for a particular class of devices all of which have an `io` ending, such as `mtio(7D)`

OPTIONS

This lists the command options with a concise summary of what each option does. The options are listed literally and in the order they appear in the SYNOPSIS section. Possible arguments to options are discussed under the option, and where appropriate, default values are supplied.

OPERANDS

This section lists the command operands and describes how they affect the actions of the command.

OUTPUT

This section describes the output - standard output, standard error, or output files - generated by the command.

RETURN VALUES

If the man page documents functions that return values, this section lists these values and describes the conditions under which they are returned. If a function can return only constant values, such as 0 or -1, these values are listed in

tagged paragraphs. Otherwise, a single paragraph describes the return values of each function. Functions declared void do not return values, so they are not discussed in RETURN VALUES.

ERRORS

On failure, most functions place an error code in the global variable `errno` indicating why they failed. This section lists alphabetically all error codes a function can generate and describes the conditions that cause each error. When more than one condition can cause the same error, each condition is described in a separate paragraph under the error code.

USAGE

This section is provided as a guidance on use. This section lists special rules, features and commands that require in-depth explanations. The subsections listed below are used to explain built-in functionality:

- Commands
- Modifiers
- Variables
- Expressions
- Input Grammar

EXAMPLES

This section provides examples of usage or of how to use a command or function. Wherever possible a complete example including command line entry and machine response is shown. Whenever an example is given, the prompt is shown as `example%` or if the user must be superuser, `example#`. Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS and USAGE sections.

ENVIRONMENT VARIABLES

This section lists any environment variables that the command or function affects, followed by a brief description of the effect.

EXIT STATUS

This section lists the values the command returns to the calling program or shell and the conditions that cause these values to be returned. Usually, zero is returned for successful completion and

values other than zero for various error conditions.

FILES

This section lists all filenames referred to by the man page, files of interest, and files created or required by commands. Each is followed by a descriptive summary or explanation.

ATTRIBUTES

This section lists characteristics of commands, utilities, and device drivers by defining the attribute type and its corresponding value. See `attributes(5)` for more information.

SEE ALSO

This section lists references to other man pages, in-house documentation and outside publications.

DIAGNOSTICS

This section lists diagnostic messages with a brief explanation of the condition causing the error.

WARNINGS

This section lists warnings about special conditions which could seriously affect your working conditions. This is not a list of diagnostics.

NOTES

This section lists additional information that does not belong anywhere else on the page. It takes the form of an aside to the user, covering points of special interest. Critical information is never covered here.

BUGS

This section describes known bugs and wherever possible, suggests workarounds.

CHAPTER

System Calls

NAME	Intro - introduction to system calls and error numbers																
SYNOPSIS	<code>#include <errno.h></code>																
DESCRIPTION	<p>This section describes all of the system calls. Most of these calls return one or more error conditions. An error condition is indicated by an otherwise impossible return value. This is almost always <code>-1</code> or the null pointer; the individual descriptions specify the details. An error number is also made available in the external variable <code>errno</code>, which is not cleared on successful calls, so it should be tested only after an error has been indicated.</p> <p>In the case of multithreaded applications, the <code>_REENTRANT</code> flag must be defined on the command line at compilation time (<code>-D_REENTRANT</code>). When the <code>_REENTRANT</code> flag is defined, <code>errno</code> becomes a macro which enables each thread to have its own <code>errno</code>. This <code>errno</code> macro can be used on either side of the assignment, just as if it were a variable.</p> <p>Applications should use bound threads rather than the <code>_lwp_*</code>() functions (see <code>thr_create(3T)</code>). Using LWPs (lightweight processes) directly is not advised because libraries are only safe to use with threads, not LWPs.</p> <p>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 defined in <code><errno.h></code>.</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">1 EPERM</td> <td>Not superuser</td> </tr> <tr> <td></td> <td>Typically this error indicates an attempt to modify a file in some way forbidden except to its owner or the super-user. It is also returned for attempts by ordinary users to do things allowed only to the super-user.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">2 ENOENT</td> <td>No such file or directory</td> </tr> <tr> <td></td> <td>A file name is specified and the file should exist but doesn't, or one of the directories in a path name does not exist.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">3 ESRCH</td> <td>No such process, LWP, or thread</td> </tr> <tr> <td></td> <td>No process can be found in the system that corresponds to the specified PID, <code>LWPID_t</code>, or <code>thread_t</code>.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">4 EINTR</td> <td>Interrupted system call</td> </tr> <tr> <td></td> <td>An asynchronous signal (such as interrupt or quit), which the user has elected to catch,</td> </tr> </table>	1 EPERM	Not superuser		Typically this error indicates an attempt to modify a file in some way forbidden except to its owner or the super-user. It is also returned for attempts by ordinary users to do things allowed only to the super-user.	2 ENOENT	No such file or directory		A file name is specified and the file should exist but doesn't, or one of the directories in a path name does not exist.	3 ESRCH	No such process, LWP, or thread		No process can be found in the system that corresponds to the specified PID, <code>LWPID_t</code> , or <code>thread_t</code> .	4 EINTR	Interrupted system call		An asynchronous signal (such as interrupt or quit), which the user has elected to catch,
1 EPERM	Not superuser																
	Typically this error indicates an attempt to modify a file in some way forbidden except to its owner or the super-user. It is also returned for attempts by ordinary users to do things allowed only to the super-user.																
2 ENOENT	No such file or directory																
	A file name is specified and the file should exist but doesn't, or one of the directories in a path name does not exist.																
3 ESRCH	No such process, LWP, or thread																
	No process can be found in the system that corresponds to the specified PID, <code>LWPID_t</code> , or <code>thread_t</code> .																
4 EINTR	Interrupted system call																
	An asynchronous signal (such as interrupt or quit), which the user has elected to catch,																

	<p>occurred during a system service function. If execution is resumed after processing the signal, it will appear as if the interrupted function call returned this error condition.</p> <p>In a multithreaded application, EINTR may be returned whenever another thread or LWP calls <code>fork(2)</code>.</p>
5 EIO	<p>I/O error</p> <p>Some physical I/O error has occurred. This error may in some cases occur on a call following the one to which it actually applies.</p>
6 ENXIO	<p>No such device or address</p> <p>I/O on a special file refers to a subdevice which does not exist, or exists beyond the limit 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.</p>
7 E2BIG	<p>Arg list too long</p> <p>An argument list longer than ARG_MAX bytes is presented to a member of the <code>exec</code> family of functions (see <code>exec(2)</code>). The argument list limit is the sum of the size of the argument list plus the size of the environment's exported shell variables.</p>
8 ENOEXEC	<p>Exec format error</p> <p>A request is made to execute a file which, although it has the appropriate permissions, does not start with a valid format (see <code>a.out(4)</code>).</p>
9 EBADF	<p>Bad file number</p> <p>Either a file descriptor refers to no open file, or a <code>read(2)</code> (respectively, <code>write(2)</code>) request is made to a file that is open only for writing (respectively, reading).</p>
10 ECHILD	<p>No child processes</p>

	<p>A <code>wait(2)</code> function was executed by a process that had no existing or unwaited-for child processes.</p>
11 EAGAIN	<p>No more processes, or no more LWPs</p> <p>For example, the <code>fork(2)</code> function failed because the system's process table is full or the user is not allowed to create any more processes, or a call failed because of insufficient memory or swap space.</p>
12 ENOMEM	<p>Not enough space</p> <p>During execution of <code>brk()</code> or <code>sbrk()</code> (see <code>brk(2)</code>), or one of the <code>exec</code> family of functions, a program asks for more space than the system is able to supply. This is not a temporary condition; the maximum size is a system parameter. On some architectures, the error may also occur if the arrangement of text, data, and stack segments requires too many segmentation registers, or if there is not enough swap space during the <code>fork(2)</code> function. If this error occurs on a resource associated with Remote File Sharing (RFS), it indicates a memory depletion which may be temporary, dependent on system activity at the time the call was invoked.</p>
13 EACCES	<p>Permission denied</p> <p>An attempt was made to access a file in a way forbidden by the protection system.</p>
14 EFAULT	<p>Bad address</p> <p>The system encountered a hardware fault in attempting to use an argument of a routine. For example, <code>errno</code> potentially may be set to <code>EFAULT</code> any time a routine that takes a pointer argument is passed an invalid address, if the system can detect the condition. Because systems will differ in their ability to reliably detect a bad address, on some implementations passing a bad address to a routine will result in undefined behavior.</p>

15 ENOTBLK	Block device required A non-block device or file was mentioned where a block device was required (for example, in a call to the <code>mount(2)</code> function).
16 EBUSY	Device busy An attempt was made to mount a device that was already mounted or an attempt was made to unmount a device on which there is an active file (open file, current directory, mounted-on file, active text segment). It will also occur if an attempt is made to enable accounting when it is already enabled. The device or resource is currently unavailable. EBUSY is also used by mutexes, semaphores, condition variables, and r/w locks, to indicate that a lock is held, and by the processor control function <code>P_ONLINE</code> .
17 EEXIST	File exists An existing file was mentioned in an inappropriate context (for example, call to the <code>link(2)</code> function).
18 EXDEV	Cross-device link A hard link to a file on another device was attempted.
19 ENODEV	No such device An attempt was made to apply an inappropriate operation to a device (for example, read a write-only device).
20 ENOTDIR	Not a directory A non-directory was specified where a directory is required (for example, in a path prefix or as an argument to the <code>chdir(2)</code> function).
21 EISDIR	Is a directory An attempt was made to write on a directory.

22 EINVAL	<p>Invalid argument</p> <p>An invalid argument was specified (for example, unmounting a non-mounted device), mentioning an undefined signal in a call to the <code>signal(3C)</code> or <code>kill(2)</code> function.</p>
23 ENFILE	<p>File table overflow</p> <p>The system file table is full (that is, <code>SYS_OPEN</code> files are open, and temporarily no more files can be opened).</p>
24 EMFILE	<p>Too many open files</p> <p>No process may have more than <code>OPEN_MAX</code> file descriptors open at a time.</p>
25 ENOTTY	<p>Inappropriate ioctl for device</p> <p>A call was made to the <code>ioctl(2)</code> function specifying a file that is not a special character device.</p>
26 ETXTBSY	<p>Text file busy (obsolete)</p> <p>An attempt was made to execute a pure-procedure program that is currently open for writing. Also an attempt to open for writing or to remove a pure-procedure program that is being executed. <i>(This message is obsolete.)</i></p>
27 EFBIG	<p>File too large</p> <p>The size of the file exceeded the limit specified by resource <code>RLIMIT_FSIZE</code>; the file size exceeds the maximum supported by the file system; or the file size exceeds the offset maximum of the file descriptor. See the <code>File Descriptor</code> subsection of the <code>DEFINITIONS</code> section below.</p>
28 ENOSPC	<p>No space left on device</p> <p>While writing an ordinary file or creating a directory entry, there is no free space left on the device. In the <code>fcntl(2)</code> function, the setting or removing of record locks on a file cannot be</p>

	accomplished because there are no more record entries left on the system.
29 EPIPE	Illegal seek A call to the <code>lseek(2)</code> function was issued to a pipe.
30 EROFS	Read-only file system An attempt to modify a file or directory was made on a device mounted read-only.
31 EMLINK	Too many links An attempt to make more than the maximum number of links, <code>LINK_MAX</code> , to a file.
32 EPIPE	Broken pipe A write on a pipe for which there is no process to read the data. This condition normally generates a signal; the error is returned if the signal is ignored.
33 EDOM	Math argument out of domain of func The argument of a function in the math package (3M) is out of the domain of the function.
34 ERANGE	Math result not representable The value of a function in the math package (3M) is not representable within machine precision.
35 ENOMSG	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 <code>msgrcv(2)</code>).
36 EIDRM	Identifier removed This error is returned to processes that resume execution due to the removal of an identifier from the file system's name space (see <code>msgctl(2)</code> , <code>semctl(2)</code> , and <code>shmctl(2)</code>).

37 ECHRNG	Channel number out of range
38 EL2NSYNC	Level 2 not synchronized
39 EL3HLT	Level 3 halted
40 EL3RST	Level 3 reset
41 ELNRNG	Link number out of range
42 EUNATCH	Protocol driver not attached
43 ENOCSI	No CSI structure available
44 EL2HLT	Level 2 halted
45 EDEADLK	Deadlock condition A deadlock situation was detected and avoided. This error pertains to file and record locking, and also applies to mutexes, semaphores, condition variables, and r/w locks.
46 ENOLCK	No record locks available There are no more locks available. The system lock table is full (see <code>fcntl(2)</code>).
47 ECANCELED	Operation canceled The associated asynchronous operation was canceled before completion.
48 ENOTSUP	Not supported This version of the system does not support this feature. Future versions of the system may provide support.
49 EDQUOT	Disc quota exceeded A <code>write(2)</code> 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.
58-59	Reserved

60 ENOSTR	Device not a stream A <code>putmsg(2)</code> or <code>getmsg(2)</code> call was attempted on a file descriptor that is not a STREAMS device.
61 ENODATA	No data available
62 ETIME	Timer expired The timer set for a STREAMS <code>ioctl(2)</code> 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 <code>ioctl()</code> operation is indeterminate. This is also returned in the case of <code>_lwp_cond_timedwait(2)</code> or <code>cond_timedwait(2)</code> .
63 ENOSR	Out of stream resources During a STREAMS <code>open(2)</code> call, either no STREAMS queues or no STREAMS head data structures were available. This is a temporary condition; one may recover from it if other processes release resources.
64 ENONET	Machine is not on the network This error is Remote File Sharing (RFS) specific. It occurs when users try to advertise, unadvertise, mount, or unmount remote resources while the machine has not done the proper startup to connect to the network.
65 ENOPKG	Package not installed This error occurs when users attempt to use a call from a package which has not been installed.
66 EREMOTE	Object is remote This error is RFS-specific. It occurs when users try to advertise a resource which is not on the local machine, or try to mount/unmount a device (or pathname) that is on a remote machine.

67 ENOLINK	<p>Link has been severed</p> <p>This error is RFS-specific. It occurs when the link (virtual circuit) connecting to a remote machine is gone.</p>
68 EADV	<p>Advertise error</p> <p>This error is RFS-specific. It occurs when users try to advertise a resource which has been advertised already, or try to stop RFS while there are resources still advertised, or try to force unmount a resource when it is still advertised.</p>
69 ESRMNT	<p>Srmount error</p> <p>This error is RFS-specific. It occurs when an attempt is made to stop RFS while resources are still mounted by remote machines, or when a resource is readvertised with a client list that does not include a remote machine that currently has the resource mounted.</p>
70 ECOMM	<p>Communication error on send</p> <p>This error is RFS-specific. It occurs when the current process is waiting for a message from a remote machine, and the virtual circuit fails.</p>
71 EPROTO	<p>Protocol error</p> <p>Some protocol error occurred. This error is device-specific, but is generally not related to a hardware failure.</p>
76 EDOTDOT	<p>Error 76</p> <p>This error is RFS-specific. A way for the server to tell the client that a process has transferred back from mount point.</p>
77 EBADMSG	<p>Not a data message</p> <p>During a <code>read(2)</code>, <code>getmsg(2)</code>, or <code>ioctl(2)</code> <code>I_RECVFD</code> call to a STREAMS device, something has come to the head of the queue that can not be processed. That something depends on the call:</p>

	read():	control information or passed file descriptor.
	getmsg():	passed file descriptor.
	ioctl():	control or data information.
78 ENAMETOOLONG	File name too long	The length of the path argument exceeds <code>PATH_MAX</code> , or the length of a path component exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect; see <code>limits(4)</code> .
79 EOVERFLOW	Value too large for defined data type.	
80 ENOTUNIQ	Name not unique on network	Given log name not unique.
81 EBADFD	File descriptor in bad state	Either a file descriptor refers to no open file or a read request was made to a file that is open only for writing.
82 EREMCHG	Remote address changed	
83 ELIBACC	Cannot access a needed share library	Trying to <code>exec</code> an <code>a.out</code> that requires a static shared library and the static shared library does not exist or the user does not have permission to use it.
84 ELIBBAD	Accessing a corrupted shared library	Trying to <code>exec</code> an <code>a.out</code> that requires a static shared library (to be linked in) and <code>exec</code> could not load the static shared library. The static shared library is probably corrupted.
85 ELIBSCN	<code>.lib</code> section in <code>a.out</code> corrupted	Trying to <code>exec</code> an <code>a.out</code> that requires a static shared library (to be linked in) and there was erroneous data in the <code>.lib</code> section of the <code>a.out</code> . The <code>.lib</code> section tells <code>exec</code> what static shared

	libraries are needed. The <code>a.out</code> is probably corrupted.
86 ELIBMAX	Attempting to link in more shared libraries than system limit Trying to <code>exec</code> an <code>a.out</code> that requires more static shared libraries than is allowed on the current configuration of the system. See <i>NFS Administration Guide</i>
87 ELIBEXEC	Cannot <code>exec</code> a shared library directly Attempting to <code>exec</code> a shared library directly.
88 EILSEQ	Error 88 Illegal byte sequence. Handle multiple characters as a single character.
89 ENOSYS	Operation not applicable
90 ELOOP	Number of symbolic links encountered during path name traversal exceeds <code>MAXSYMLINKS</code>
91 ESTART	Restartable system call Interrupted system call should be restarted.
92 ESTRPIPE	If pipe/FIFO, don't sleep in stream head Streams pipe error (not externally visible).
93 ENOTEMPTY	Directory not empty
94 EUSERS	Too many users
95 ENOTSOCK	Socket operation on non-socket
96 EDESTADDRREQ	Destination address required A required address was omitted from an operation on a transport endpoint. Destination address required.
97 EMGSIZE	Message too long

	A message sent on a transport provider was larger than the internal message buffer or some other network limit.
98 EPROTOTYPE	Protocol wrong type for socket A protocol was specified that does not support the semantics of the socket type requested.
99 ENOPROTOOPT	Protocol not available A bad option or level was specified when getting or setting options for a protocol.
120 EPROTONOSUPPORT	Protocol not supported The protocol has not been configured into the system or no implementation for it exists.
121 ESOCKTNOSUPPORT	Socket type not supported The support for the socket type has not been configured into the system or no implementation for it exists.
122 EOPNOTSUPP	Operation not supported on transport endpoint For example, trying to accept a connection on a datagram transport endpoint.
123 EPFNOSUPPORT	Protocol family not supported The protocol family has not been configured into the system or no implementation for it exists. Used for the Internet protocols.
124 EAFNOSUPPORT	Address family not supported by protocol family An address incompatible with the requested protocol was used.
125 EADDRINUSE	Address already in use User attempted to use an address already in use, and the protocol does not allow this.
126 EADDRNOTAVAIL	Cannot assign requested address

	Results from an attempt to create a transport endpoint with an address not on the current machine.
127 ENETDOWN	Network is down Operation encountered a dead network.
128 ENETUNREACH	Network is unreachable Operation was attempted to an unreachable network.
129 ENETRESET	Network dropped connection because of reset The host you were connected to crashed and rebooted.
130 ECONNABORTED	Software caused connection abort A connection abort was caused internal to your host machine.
131 ECONNRESET	Connection reset by peer A connection was forcibly closed by a peer. This normally results from a loss of the connection on the remote host due to a timeout or a reboot.
132 ENOBUFS	No buffer space available An operation on a transport endpoint or pipe was not performed because the system lacked sufficient buffer space or because a queue was full.
133 EISCONN	Transport endpoint is already connected A connect request was made on an already connected transport endpoint; or, a <code>sendto(3N)</code> or <code>sendmsg(3N)</code> request on a connected transport endpoint specified a destination when already connected.
134 ENOTCONN	Transport endpoint is not connected A request to send or receive data was disallowed because the transport endpoint is not connected

	and (when sending a datagram) no address was supplied.
143 ESHUTDOWN	Cannot send after transport endpoint shutdown A request to send data was disallowed because the transport endpoint has already been shut down.
144 ETOOMANYREFS	Too many references: cannot splice
145 ETIMEDOUT	Connection timed out A <code>connect(3N)</code> or <code>send(3N)</code> request failed because the connected party did not properly respond after a period of time; or a <code>write(2)</code> or <code>fsync(3C)</code> request failed because a file is on an NFS file system mounted with the <i>soft</i> option.
146 ECONNREFUSED	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 remote host.
147 EHOSTDOWN	Host is down A transport provider operation failed because the destination host was down.
148 EHOSTUNREACH	No route to host A transport provider operation was attempted to an unreachable host.
149 EALREADY	Operation already in progress An operation was attempted on a non-blocking object that already had an operation in progress.
150 EINPROGRESS	Operation now in progress An operation that takes a long time to complete (such as a <code>connect()</code>) was attempted on a non-blocking object.
151 ESTALE	Stale NFS file handle

DEFINITIONS

Background Process Group	Any process group that is not the foreground process group of a session that has established a connection with a controlling terminal.
Controlling Process	A session leader that established a connection to a controlling terminal.
Controlling Terminal	A terminal that is associated with a session. Each session may have, at most, one controlling terminal associated with it and a controlling terminal may be associated with only one session. Certain input sequences from the controlling terminal cause signals to be sent to process groups in the session associated with the controlling terminal; see <code>termio(7I)</code> .
Directory	Directories organize files into a hierarchical system where directories are the nodes in the hierarchy. A directory is a file that catalogs the list of files, including directories (sub-directories), that are directly beneath it in the hierarchy. Entries in a directory file are called links. A link associates a file identifier with a filename. By convention, a directory contains at least two links, <code>.</code> (dot) and <code>..</code> (dot-dot). The link called dot refers to the directory itself while dot-dot refers to its parent directory. The root directory, which is the top-most node of the hierarchy, has itself as its parent directory. The pathname of the root directory is <code>/</code> and the parent directory of the root directory is <code>/</code> .
Downstream	In a stream, the direction from stream head to driver.
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 log driver (see <code>log(7D)</code>), which is not associated with a hardware device.
Effective User ID and Effective Group ID	An active process has an effective user ID and an effective group ID that are used to determine file access permissions (see below). The effective user ID and effective group ID are equal to the process's real user ID and real group ID, respectively, unless the process or one of its ancestors evolved from a file that had the set-user-ID bit or set-group-ID bit set (see <code>exec(2)</code>).
File Access Permissions	Read, write, and execute/search permissions for a file are granted to a process if one or more of the following are true: <ul style="list-style-type: none"> ■ The effective user ID of the process is super-user. ■ The effective user ID of the process matches the user ID of the owner of the file and the appropriate access bit of the "owner" portion (0700) of the file mode is set. ■ The effective user ID of the process does not match the user ID of the owner of the file, but either the effective group ID or one of the supplementary

	<p>group IDs of the process match the group ID of the file and the appropriate access bit of the “group” portion (0070) of the file mode is set.</p> <ul style="list-style-type: none"> ■ The effective user ID of the process does not match the user ID of the owner of the file, and neither the effective group ID nor any of the supplementary group IDs of the process match the group ID of the file, but the appropriate access bit of the “other” portion (0007) of the file mode is set. <p>Otherwise, the corresponding permissions are denied.</p>
File Descriptor	<p>A file descriptor is a small integer used to perform I/O on a file. The value of a file descriptor is from 0 to (NOFILES-1). A process may have no more than NOFILES file descriptors open simultaneously. A file descriptor is returned by calls such as <code>open(2)</code> or <code>pipe(2)</code>. The file descriptor is used as an argument by calls such as <code>read(2)</code>, <code>write(2)</code>, <code>ioctl(2)</code>, and <code>close(2)</code>.</p> <p>Each file descriptor has a corresponding offset maximum. For regular files that were opened without setting the <code>O_LARGEFILE</code> flag, the offset maximum is 2 Gbyte - 1 byte ($2^{31} - 1$ bytes). For regular files that were opened with the <code>O_LARGEFILE</code> flag set, the offset maximum is $2^{63} - 1$ bytes.</p>
File Name	<p>Names consisting of 1 to <code>NAME_MAX</code> characters may be used to name an ordinary file, special file or directory.</p> <p>These characters may be selected from the set of all character values excluding <code>\0</code> (null) and the ASCII code for <code>/</code> (slash).</p> <p>Note that it is generally unwise to use <code>*</code>, <code>?</code>, <code>[</code>, or <code>]</code> as part of file names because of the special meaning attached to these characters by the shell (see <code>sh(1)</code>, <code>csh(1)</code>, and <code>ksh(1)</code>). Although permitted, the use of unprintable characters in file names should be avoided.</p> <p>A file name is sometimes referred to as a pathname component. The interpretation of a pathname component is dependent on the values of <code>NAME_MAX</code> and <code>_POSIX_NO_TRUNC</code> associated with the path prefix of that component. If any pathname component is longer than <code>NAME_MAX</code> and <code>_POSIX_NO_TRUNC</code> is in effect for the path prefix of that component (see <code>fpathconf(2)</code> and <code>limits(4)</code>), it shall be considered an error condition in that implementation. Otherwise, the implementation shall use the first <code>NAME_MAX</code> bytes of the pathname component.</p>
Foreground Process Group	<p>Each session that has established a connection with a controlling terminal will distinguish one process group of the session as the foreground process group of the controlling terminal. This group has certain privileges when accessing its controlling terminal that are denied to background process groups.</p>
{IOV_MAX}	<p>Maximum number of entries in a <code>struct iovec</code> array.</p>

{LIMIT}	The braces notation, {LIMIT}, is used to denote a magnitude limitation imposed by the implementation. This indicates a value which may be defined by a header file (without the braces), or the actual value may be obtained at runtime by a call to the configuration inquiry <code>pathconf(2)</code> with the name argument <code>_PC_LIMIT</code> .
Masks	The file mode creation mask of the process used during any create function calls to turn off permission bits in the <i>mode</i> argument supplied. Bit positions that are set in <code>umask(cmask)</code> are cleared in the mode of the created file.
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.
Message Queue Identifier	<p>A message queue identifier (<code>msqid</code>) is a unique positive integer created by a <code>msgget(2)</code> call. Each <code>msqid</code> has a message queue and a data structure associated with it. The data structure is referred to as <code>msqid_ds</code> and contains the following members:</p> <pre> struct ipc_perm msg_perm; struct msg *msg_first; struct msg *msg_last; ulong_t msg_cbytes; ulong_t msg_qnum; ulong_t msg_qbytes; pid_t msg_lspid; pid_t msg_lrpid; time_t msg_stime; time_t msg_rtime; time_t msg_ctime; </pre> <p>The following are descriptions of the <code>msqid_ds</code> structure members:</p> <p>The <code>msg_perm</code> member is an <code>ipc_perm</code> structure that specifies the message operation permission (see below). This structure includes the following members:</p> <pre> uid_t cuid; /* creator user id */ gid_t cgid; /* creator group id */ uid_t uid; /* user id */ </pre>

```

gid_t   gid;    /* group id */
mode_t  mode;   /* r/w permission */
ulong_t seq;    /* slot usage sequence # */
key_t   key;    /* key */

```

The `*msg_first` member is a pointer to the first message on the queue.

The `*msg_last` member is a pointer to the last message on the queue.

The `msg_cbytes` member is the current number of bytes on the queue.

The `msg_qnum` member is the number of messages currently on the queue.

The `msg_qbytes` member is the maximum number of bytes allowed on the queue.

The `msg_lspid` member is the process ID of the last process that performed a `msgsnd()` operation.

The `msg_lrpid` member is the process id of the last process that performed a `msgrcv()` operation.

The `msg_stime` member is the time of the last `msgsnd()` operation.

The `msg_rtime` member is the time of the last `msgrcv()` operation.

The `msg_ctime` member is the time of the last `msgctl()` operation that changed a member of the above structure.

Message Operation Permissions

In the `msgctl(2)`, `msgget(2)`, `msgrcv(2)`, and `msgsnd(2)` function 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
00040  READ by group
00020  WRITE by group
00004  READ by others
00002  WRITE by others

```

Read and write permissions for a `msgid` 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.cuid` or `msg_perm.uid` in the data structure associated with `msgid` and the appropriate bit of the “user” portion (0600) of `msg_perm.mode` is set.

- Any group ID in the process credentials from the set (`cr_gid`, `cr_groups`) matches `msg_perm.cgid` or `msg_perm.gid` and the appropriate bit of the “group” portion (060) of `msg_perm.mode` is set.
- The appropriate bit of the “other” portion (006) of `msg_perm.mode` is set.”

Otherwise, the corresponding permissions are denied.

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.
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.
Offset Maximum	An offset maximum is an attribute of an open file description representing the largest value that can be used as a file offset.
Orphaned Process Group	A process group in which the parent of every member in the group is either itself a member of the group, or is not a member of the process group’s session.
Path Name	<p>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.</p> <p>If a path name begins with a slash, the path search begins at the root directory. Otherwise, the search begins from the current working directory.</p> <p>A slash by itself names the root directory.</p> <p>Unless specifically stated otherwise, the null path name is treated as if it named a non-existent file.</p>
Process ID	Each process in the system is uniquely identified during its lifetime by a positive integer called a process ID. A process ID may not be reused by the system until the process lifetime, process group lifetime, and session lifetime ends for any process ID, process group ID, and session ID equal to that process ID. Within a process, there are threads with thread id’s, called <code>thread_t</code> and <code>LWPID_t</code> . These threads are not visible to the outside process.

Parent Process ID	A new process is created by a currently active process (see <code>fork(2)</code>). The parent process ID of a process is the process ID of its creator.
Privilege	Having appropriate privilege means having the capability to override system restrictions.
Process Group	Each process in the system is a member of a process group that is identified by a process group ID. Any process that is not a process group leader may create a new process group and become its leader. Any process that is not a process group leader may join an existing process group that shares the same session as the process. A newly created process joins the process group of its parent.
Process Group Leader	A process group leader is a process whose process ID is the same as its process group ID.
Process Group ID	Each active process is a member of a process group and 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 <code>kill(2)</code>).
Process Lifetime	A process lifetime begins when the process is forked and ends after it exits, when its termination has been acknowledged by its parent process. See <code>wait(2)</code> .
Process Group Lifetime	A process group lifetime begins when the process group is created by its process group leader, and ends when the lifetime of the last process in the group ends or when the last process in the group leaves the group.
Processor Set ID	The processors in a system may be divided into subsets, known as processor sets. A process bound to one of these sets will run only on processors in that set, and the processors in the set will normally run only processes that have been bound to the set. Each active processor set is identified by a positive integer. See <code>pset_create(2)</code> .
Read Queue	In a stream, the message queue in a module or driver containing messages moving upstream.
Real User ID and Real Group ID	Each user allowed on the system is identified by a positive integer (0 to <code>MAXUID</code>) called a real user ID. Each user is also a member of a group. The group is identified by a positive integer called the real group ID. An active process has a real user ID and real group ID that are set to the real user ID and real group ID, respectively, of the user responsible for the creation of the process.

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 of a process need not be the root directory of the root file system.
Saved Resource Limits	Saved resource limits is an attribute of a process that provides some flexibility in the handling of unrepresentable resource limits, as described in the <code>exec</code> family of functions and <code>setrlimit(2)</code> .
Saved User ID and Saved Group ID	The saved user ID and saved group ID are the values of the effective user ID and effective group ID prior to an <code>exec</code> of a file whose set user or set group file mode bit has been set (see <code>exec(2)</code>).
Semaphore Identifier	<p>A semaphore identifier (<code>semid</code>) is a unique positive integer created by a <code>semget(2)</code> call. Each <code>semid</code> has a set of semaphores and a data structure associated with it. The data structure is referred to as <code>semid_ds</code> and contains the following members:</p> <pre> struct ipc_perm sem_perm; /* operation permission struct */ struct sem *sem_base; /* ptr to first semaphore in set */ ushort_t 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 */ </pre> <p>The following are descriptions of the <code>semid_ds</code> structure members:</p> <p>The <code>sem_perm</code> member is an <code>ipc_perm</code> structure that specifies the semaphore operation permission (see below). This structure includes the following members:</p> <pre> uid_t uid; /* user id */ gid_t gid; /* group id */ uid_t cuid; /* creator user id */ gid_t cgid; /* creator group id */ mode_t mode; /* r/a permission */ ulong_t seq; /* slot usage sequence number */ key_t key; /* key */ </pre> <p>The <code>sem_nsems</code> member is equal to the number of semaphores in the set. Each semaphore in the set is referenced by a nonnegative integer referred to as a <code>sem_num</code>. <code>sem_num</code> values run sequentially from 0 to the value of <code>sem_nsems</code> minus 1.</p> <p>The <code>sem_otime</code> member is the time of the last <code>semop(2)</code> operation.</p>

The `sem_ctime` member is the time of the last `semctl(2)` operation that changed a member of the above structure.

A semaphore is a data structure called `sem` that contains the following members:

```
ushort_t  semval;    /* semaphore value */
pid_t     sempid;   /* pid of last operation */
ushort_t  semncnt;  /* # awaiting semval > cval */
ushort_t  semzcnt;  /* # awaiting semval = 0 */
```

The following are descriptions of the `sem` structure members:

The `semval` member is a non-negative integer that is the actual value of the semaphore.

The `sempid` member is equal to the process ID of the last process that performed a semaphore operation on this semaphore.

The `semncnt` member is a count of the number of processes that are currently suspended awaiting this semaphore's `semval` to become greater than its current value.

The `semzcnt` member is a count of the number of processes that are currently suspended awaiting this semaphore's `semval` to become 0.

Semaphore Operation Permissions

In the `semop(2)` and `semctl(2)` function 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
00040 READ by group
00020 ALTER by group
00004 READ by others
00002 ALTER by others
```

Read and alter permissions for 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.cuid` or `sem_perm.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 group ID of the process matches `sem_perm.cgid` or `sem_perm.gid` and the appropriate bit of the “group” portion (060) of `sem_perm.mode` is set.
- The appropriate bit of the “other” portion (06) of `sem_perm.mode` is set.

Otherwise, the corresponding permissions are denied.

Session A session is a group of processes identified by a common ID called a session ID, capable of establishing a connection with a controlling terminal. Any process that is not a process group leader may create a new session and process group, becoming the session leader of the session and process group leader of the process group. A newly created process joins the session of its creator.

Session ID Each session in the system is uniquely identified during its lifetime by a positive integer called a session ID, the process ID of its session leader.

Session Leader A session leader is a process whose session ID is the same as its process and process group ID.

Session Lifetime A session lifetime begins when the session is created by its session leader, and ends when the lifetime of the last process that is a member of the session ends, or when the last process that is a member in the session leaves the session.

Shared Memory Identifier A shared memory identifier (`shmid`) is a unique positive integer created by a `semget(2)` call. Each `shmid` has a segment of memory (referred to as a shared memory segment) and a data structure associated with it. (Note that these shared memory segments must be explicitly removed by the user after the last reference to them is removed.) 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 */
struct region    *shm_reg;     /* ptr to region structure */
char             pad[4];       /* for swap compatibility */
pid_t            shm_lpid;     /* pid of last operation */
pid_t            shm_cpid;     /* creator pid */
ushort_t         shm_nattch;   /* number of current attaches */
ushort_t         shm_cnattch;  /* used only for shminfo */
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 */

```

The following are descriptions of the `shmid_ds` structure members:

The `shm_perm` member is an `ipc_perm` structure that specifies the shared memory operation permission (see below). This structure includes the following members:

```

uid_t      cuid; /* creator user id */
gid_t      cgid; /* creator group id */
uid_t      uid;  /* user id */
gid_t      gid;  /* group id */
mode_t     mode; /* r/w permission */
ulong_t    seq;  /* slot usage sequence # */
key_t      key;  /* key */

```

The `shm_segsz` member specifies the size of the shared memory segment in bytes.

The `shm_cpid` member is the process ID of the process that created the shared memory identifier.

The `shm_lpid` member is the process ID of the last process that performed a `shmat()` or `shmdt()` operation (see `shmop(2)`).

The `shm_nattch` member is the number of processes that currently have this segment attached.

The `shm_atime` member is the time of the last `shmat()` operation (see `shmop(2)`).

The `shm_dtime` member is the time of the last `shmdt()` operation (see `shmop(2)`).

The `shm_ctime` member is the time of the last `shmctl(2)` operation that changed one of the members of the above structure.

Shared Memory Operation Permissions

In the `shmctl(2)`, `shmat()`, and `shmdt()` (see `shmop(2)`) function 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
00040  READ by group
00020  WRITE by group
00004  READ by others
00002  WRITE by others

```

Read and write permissions for 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.cuid` or `shm_perm.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 group ID of the process matches `shm_perm.cgid` or `shm_perm.gid` and the appropriate bit of the “group” portion (060) of `shm_perm.mode` is set.
- The appropriate bit of the “other” portion (06) of `shm_perm.mode` is set.

Otherwise, the corresponding permissions are denied.

Special Processes

The process with ID 0 and the process with ID 1 are special processes referred to as `proc0` and `proc1`; see `kill(2)`. `proc0` is the process scheduler. `proc1` is the initialization process (*init*); `proc1` is the ancestor of every other process in the system and is used to control the process structure.

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 principal functions of the stream head are processing STREAMS-related system calls and passing data and information between a user process and the stream.

Super-user	A process is recognized as a super-user process and is granted special privileges, such as immunity from file permissions, if its effective user ID is 0.
Upstream	In a stream, the direction from driver to stream head.
Write Queue	In a stream, the message queue in a module or driver containing messages moving downstream.

LIST OF FUNCTIONS

Name	Description
Intro(2)	introduction to system calls and error numbers
__sparc_utrap_install(2)	install a SPARC V9 user trap handler
_exit(2)	See exit(2)
_lwp_cond_broadcast(2)	See _lwp_cond_signal(2)
_lwp_cond_signal(2)	signal a condition variable
_lwp_cond_timedwait(2)	See _lwp_cond_wait(2)
_lwp_cond_wait(2)	wait on a condition variable
_lwp_continue(2)	See _lwp_suspend(2)
_lwp_create(2)	create a new light-weight process
_lwp_exit(2)	terminate the calling LWP
_lwp_getprivate(2)	See _lwp_setprivate(2)
_lwp_info(2)	return the time-accounting information of a single LWP
_lwp_kill(2)	send a signal to a LWP
_lwp_makecontext(2)	initialize an LWP context
_lwp_mutex_lock(2)	mutual exclusion
_lwp_mutex_trylock(2)	See _lwp_mutex_lock(2)
_lwp_mutex_unlock(2)	See _lwp_mutex_lock(2)
_lwp_self(2)	get LWP identifier

<code>_lwp_sema_init(2)</code>	See <code>_lwp_sema_wait(2)</code>
<code>_lwp_sema_post(2)</code>	See <code>_lwp_sema_wait(2)</code>
<code>_lwp_sema_trywait(2)</code>	See <code>_lwp_sema_wait(2)</code>
<code>_lwp_sema_wait(2)</code>	semaphore operations
<code>_lwp_setprivate(2)</code>	set or
<code>_lwp_sigredirect(2)</code>	See <code>_signotifywait(2)</code>
<code>_lwp_suspend(2)</code>	continue or suspend LWP execution
<code>_lwp_wait(2)</code>	wait for a LWP to terminate
<code>_signotifywait(2)</code>	deliver process signals to specific LWPs
<code>access(2)</code>	determine accessibility of a file
<code>acct(2)</code>	enable or disable process accounting
<code>acl(2)</code>	get or set a file's Access Control List (ACL)
<code>adjtime(2)</code>	correct the time to allow synchronization of the system clock
<code>alarm(2)</code>	set a process alarm clock
<code>audit(2)</code>	write a record to the audit log
<code>auditon(2)</code>	manipulate auditing
<code>auditsvc(2)</code>	write audit log to specified file descriptor
<code>brk(2)</code>	change the amount of space allocated for the calling process's data segment
<code>chdir(2)</code>	change working directory
<code>chmod(2)</code>	change access permission mode of file
<code>chown(2)</code>	change owner and group of a file

chroot(2)	change root directory
close(2)	close a file descriptor
creat(2)	create a new file or rewrite an existing one
dup(2)	duplicate an open file descriptor
exec(2)	execute a file
execl(2)	See exec(2)
execle(2)	See exec(2)
execlp(2)	See exec(2)
execv(2)	See exec(2)
execve(2)	See exec(2)
execvp(2)	See exec(2)
exit(2)	terminate process
fac1(2)	See ac1(2)
fchdir(2)	See chdir(2)
fchmod(2)	See chmod(2)
fchown(2)	See chown(2)
fchroot(2)	See chroot(2)
fcntl(2)	file control
fork(2)	create a new process
fork1(2)	See fork(2)
fpathconf(2)	get configurable pathname variables
fstat(2)	See stat(2)
fstatvfs(2)	See statvfs(2)
getaudit(2)	get and set process audit information

<code>getaudit(2)</code>	get and set user audit identity
<code>getcontext(2)</code>	get and set current user context
<code>getdents(2)</code>	read directory entries and put in a file system independent format
<code>getegid(2)</code>	See <code>getuid(2)</code>
<code>geteuid(2)</code>	See <code>getuid(2)</code>
<code>getgid(2)</code>	See <code>getuid(2)</code>
<code>getgroups(2)</code>	get or set supplementary group access list IDs
<code>getitimer(2)</code>	get or set value of interval timer
<code>getmsg(2)</code>	get next message off a stream
<code>getpgid(2)</code>	See <code>getpid(2)</code>
<code>getpgrp(2)</code>	See <code>getpid(2)</code>
<code>getpid(2)</code>	get process, process group, and parent process IDs
<code>getpmsg(2)</code>	See <code>getmsg(2)</code>
<code>getppid(2)</code>	See <code>getpid(2)</code>
<code>getrlimit(2)</code>	control maximum system resource consumption
<code>getsid(2)</code>	get process group ID of session leader
<code>getuid(2)</code>	get real user, effective user, real group, and effective group IDs
<code>intro(2)</code>	See <code>Intro(2)</code>
<code>ioctl(2)</code>	control device
<code>kill(2)</code>	send a signal to a process or a group of processes
<code>lchown(2)</code>	See <code>chown(2)</code>

<code>link(2)</code>	link to a file
<code>llseek(2)</code>	move extended read/write file pointer
<code>lseek(2)</code>	move read/write file pointer
<code>lstat(2)</code>	See <code>stat(2)</code>
<code>memcntl(2)</code>	memory management control
<code>mincore(2)</code>	determine residency of memory pages
<code>mkdir(2)</code>	make a directory
<code>mknod(2)</code>	make a directory, or a special or ordinary file
<code>mmap(2)</code>	map pages of memory
<code>mount(2)</code>	mount a file system
<code>mprotect(2)</code>	set protection of memory mapping
<code>msgctl(2)</code>	message control operations
<code>msgget(2)</code>	get message queue
<code>msgrcv(2)</code>	message receive operation
<code>msgsnd(2)</code>	message send operation
<code>munmap(2)</code>	unmap pages of memory
<code>nice(2)</code>	change priority of a process
<code>ntp_adjtime(2)</code>	adjust local clock parameters
<code>ntp_gettime(2)</code>	get local clock values
<code>open(2)</code>	open a file
<code>p_online(2)</code>	return or change processor operational status
<code>pathconf(2)</code>	See <code>fpathconf(2)</code>

<code>pause(2)</code>	suspend process until signal
<code>pcsample(2)</code>	program execution time profile
<code>pipe(2)</code>	create an interprocess channel
<code>poll(2)</code>	input/output multiplexing
<code>pread(2)</code>	See <code>read(2)</code>
<code>prcntl(2)</code>	process scheduler control
<code>prcntlset(2)</code>	generalized process scheduler control
<code>processor_bind(2)</code>	bind LWPs to a processor
<code>processor_info(2)</code>	determine type and status of a processor
<code>profil(2)</code>	execution time profile
<code>pset_assign(2)</code>	See <code>pset_create(2)</code>
<code>pset_bind(2)</code>	bind LWPs to a set of processors
<code>pset_create(2)</code>	manage sets of processors
<code>pset_destroy(2)</code>	See <code>pset_create(2)</code>
<code>pset_info(2)</code>	get information about a processor set
<code>ptrace(2)</code>	allows a parent process to control the execution of a child process
<code>putmsg(2)</code>	send a message on a stream
<code>putpmsg(2)</code>	See <code>putmsg(2)</code>
<code>pwrite(2)</code>	See <code>write(2)</code>
<code>read(2)</code>	read from file
<code>readlink(2)</code>	read the contents of a symbolic link
<code>readv(2)</code>	See <code>read(2)</code>
<code>rename(2)</code>	change the name of a file

<code>resolvepath(2)</code>	resolve all symbolic links of a path name
<code>rmdir(2)</code>	remove a directory
<code>sbrk(2)</code>	See <code>brk(2)</code>
<code>semctl(2)</code>	semaphore control operations
<code>semget(2)</code>	get set of semaphores
<code>semop(2)</code>	semaphore operations
<code>setaudit(2)</code>	See <code>getaudit(2)</code>
<code>setaudit(2)</code>	See <code>getaudit(2)</code>
<code>setcontext(2)</code>	See <code>getcontext(2)</code>
<code>setegid(2)</code>	See <code>setuid(2)</code>
<code>seteuid(2)</code>	See <code>setuid(2)</code>
<code>setgid(2)</code>	See <code>setuid(2)</code>
<code>setgroups(2)</code>	See <code>getgroups(2)</code>
<code>setitimer(2)</code>	See <code>getitimer(2)</code>
<code>setpgid(2)</code>	set process group ID
<code>setpgrp(2)</code>	set process group ID
<code>setregid(2)</code>	set real and effective group IDs
<code>setreuid(2)</code>	set real and effective user IDs
<code>setrlimit(2)</code>	See <code>getrlimit(2)</code>
<code>setsid(2)</code>	create session and set process group ID
<code>setuid(2)</code>	set user and group IDs
<code>shmat(2)</code>	See <code>shmop(2)</code>
<code>shmctl(2)</code>	shared memory control operations
<code>shmdt(2)</code>	See <code>shmop(2)</code>

<code>shmget(2)</code>	get shared memory segment identifier
<code>shmop(2)</code>	shared memory operations
<code>sigaction(2)</code>	detailed signal management
<code>sigaltstack(2)</code>	set or get signal alternate stack context
<code>sigpending(2)</code>	examine signals that are blocked and pending
<code>sigprocmask(2)</code>	change and/or examine caller's signal mask
<code>sigsend(2)</code>	send a signal to a process or a group of processes
<code>sigsendset(2)</code>	See <code>sigsend(2)</code>
<code>sigsuspend(2)</code>	install a signal mask and suspend caller until signal
<code>sigwait(2)</code>	wait until a signal is posted
<code>stat(2)</code>	get file status
<code>statvfs(2)</code>	get file system information
<code>stime(2)</code>	set system time and date
<code>swapctl(2)</code>	manage swap space
<code>symlink(2)</code>	make a symbolic link to a file
<code>sync(2)</code>	update super block
<code>sysfs(2)</code>	get file system type information
<code>sysinfo(2)</code>	get and set system information strings
<code>time(2)</code>	get time
<code>times(2)</code>	get process and child process times

<code>uadmin(2)</code>	administrative control
<code>ulimit(2)</code>	get and set process limits
<code>umask(2)</code>	set and get file creation mask
<code>umount(2)</code>	unmount a file system
<code>uname(2)</code>	get name of current operating system
<code>unlink(2)</code>	remove directory entry
<code>ustat(2)</code>	get file system statistics
<code>utime(2)</code>	set file access and modification times
<code>utimes(2)</code>	set file access and modification times
<code>vfork(2)</code>	spawn new process in a virtual memory efficient way
<code>vhangup(2)</code>	virtually hangup the current controlling terminal
<code>wait(2)</code>	wait for child process to stop or terminate
<code>waitid(2)</code>	wait for child process to change state
<code>waitpid(2)</code>	wait for child process to change state
<code>write(2)</code>	write on a file
<code>writev(2)</code>	See <code>write(2)</code>
<code>yield(2)</code>	yield execution to another lightweight process

NAME	access – determine accessibility of a file
SYNOPSIS	<pre>#include <unistd.h> int access(const char *path, int amode);</pre>
DESCRIPTION	<p>The access() function checks the file named by the pathname pointed to by the <i>path</i> argument for accessibility according to the bit pattern contained in <i>amode</i>, using the real user ID in place of the effective user ID and the real group ID in place of the effective group ID. This allows a setuid process to verify that the user running it would have had permission to access this file.</p> <p>The value of <i>amode</i> is either the bitwise inclusive OR of the access permissions to be checked (R_OK, W_OK, X_OK) or the existence test, F_OK.</p> <p>These constants are defined in <unistd.h> as follows:</p> <p>R_OK Test for read permission.</p> <p>W_OK Test for write permission.</p> <p>X_OK Test for execute or search permission.</p> <p>F_OK Check existence of file</p> <p>See intro(2) for additional information about "File Access Permission".</p> <p>If any access permissions are to be checked, each will be checked individually, as described in intro(2). If the process has appropriate privileges, an implementation may indicate success for X_OK even if none of the execute file permission bits are set.</p>
RETURN VALUES	If the requested access is permitted, access() succeeds and returns 0. Otherwise, -1 is returned and errno is set to indicate the error.
ERRORS	<p>The access() function will fail if:</p> <p>EACCES Permission bits of the file mode do not permit the requested access, or search permission is denied on a component of the path prefix.</p> <p>EFAULT <i>path</i> points to an illegal address.</p> <p>EINTR A signal was caught during the access() function.</p> <p>ELOOP Too many symbolic links were encountered in resolving <i>path</i>.</p> <p>ENAMETOOLONG The length of the <i>path</i> argument exceeds PATH_MAX, or a pathname component is longer</p>

than `NAME_MAX` while `_POSIX_NO_TRUNC` is in effect.

ENOENT

A component of *path* does not name an existing file or *path* is an empty string.

ENOLINK

path points to a remote machine and the link to that machine is no longer active.

ENOTDIR

A component of the path prefix is not a directory.

EROFS

Write access is requested for a file on a read-only file system.

The `access()` function may fail if:

EINVAL

The value of the *amode* argument is invalid.

ENAMETOOLONG

Pathname resolution of a symbolic link produced an intermediate result whose length exceeds `PATH_MAX`.

ETXTBSY

Write access is requested for a pure procedure (shared text) file that is being executed.

USAGE

Additional values of *amode* other than the set defined in the description may be valid, for example, if a system has extended access controls.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

`intro(2)`, `chmod(2)`, `stat(2)`, `attributes(5)`

NAME	acct – enable or disable process accounting
SYNOPSIS	<pre>#include <unistd.h> int acct(const char *path);</pre>
DESCRIPTION	<p>The acct() function enables or disables the system process accounting routine. If the routine is enabled, an accounting record will be written in an accounting file for each process that terminates. The termination of a process can be caused by either an exit(2) call or a signal(3C). The effective user ID of the process calling acct() must be super-user.</p> <p>The <i>path</i> argument points to the pathname of the accounting file, whose file format is described on the acct(4) manual page.</p> <p>The accounting routine is enabled if <i>path</i> is non-zero and no errors occur during the function. It is disabled if <i>path</i> is <code>(char *)NULL</code> and no errors occur during the function.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The acct() function will fail if:</p> <p>EACCES The file named by <i>path</i> is not an ordinary file.</p> <p>EBUSY An attempt is being made to enable accounting using the same file that is currently being used.</p> <p>EFAULT The <i>path</i> argument points to an illegal address.</p> <p>ELOOP Too many symbolic links were encountered in translating <i>path</i>.</p> <p>ENAMETOOLONG The length of the <i>path</i> argument exceeds <code>PATH_MAX</code>, or the length of a <i>path</i> argument exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.</p> <p>ENOENT One or more components of the accounting file pathname do not exist.</p> <p>ENOTDIR A component of the path prefix is not a directory.</p> <p>EPERM The effective user of the calling process is not super-user.</p> <p>EROFS The named file resides on a read-only file system.</p>

acct(2)

System Calls

SEE ALSO | `exit(2)`, `signal(3C)`, `acct(4)`

NAME	acl, facl – get or set a file’s Access Control List (ACL)
SYNOPSIS	<pre>#include <sys/acl.h> int acl(char * <i>pathp</i>, int <i>cmd</i>, int <i>nentries</i>, aclent_t * <i>aclbufp</i>); int facl(int <i>fildes</i>, int <i>cmd</i>, int <i>nentries</i>, aclent_t * <i>aclbufp</i>);</pre>
DESCRIPTION	<p>The acl() and facl() functions get or set the ACL of a file whose name is given by <i>pathp</i> or referenced by the open file descriptor <i>fildes</i> . The <i>nentries</i> argument specifies how many ACL entries fit into buffer <i>aclbufp</i> . The acl() function is used to manipulate ACL on file system objects.</p> <p>The following values for <i>cmd</i> are supported:</p> <p>SETACL <i>nentries</i> ACL entries, specified in buffer <i>aclbufp</i> , are stored in the file’s ACL. This command can only be executed by a process that has an effective user ID equal to the owner of the file. All directories in the path name must be searchable.</p> <p>GETACL Buffer <i>aclbufp</i> is filled with the file’s ACL entries. Read access to the file is not required, but all directories in the path name must be searchable.</p> <p>GETACL CNT The number of entries in the file’s ACL is returned. Read access to the file is not required, but all directories in the path name must be searchable.</p>
RETURN VALUES	Upon successful completion, acl() and facl() return 0 if <i>cmd</i> is SETACL . If <i>cmd</i> is GETACL or GETACL CNT , the number of ACL entries is returned. Otherwise, -1 is returned and <i>errno</i> is set to indicate the error.
ERRORS	<p>The acl() function will fail if:</p> <p>EACCESS The caller does not have access to a component of the pathname.</p> <p>EFAULT The <i>pathp</i> or <i>aclbufp</i> argument points to an illegal address.</p> <p>EINVAL The <i>cmd</i> argument is not GETACL , SETACL , or GETACL CNT ; the <i>cmd</i> argument is SETACL and <i>nentries</i> is less than 3; or the <i>cmd</i> argument is SETACL and the ACL specified in <i>aclbufp</i> is not valid.</p> <p>EIO A disk I/O error has occurred while storing or retrieving the ACL.</p> <p>ENOENT A component of the path does not exist.</p>

- ENOSPC** The *cmd* argument is `GETACL` and *nentries* is less than the number of entries in the file's ACL , or the *cmd* argument is `SETACL` and there is insufficient space in the file system to store the ACL .
- ENOTDIR** A component of the path specified by *pathp* is not a directory, or the *cmd* argument is `SETACL` and an attempt is made to set a default ACL on a file type other than a directory.
- ENOSYS** The *cmd* argument is `SETACL` and the file specified by *pathp* resides on a file system that does not support ACLs , or the `acl()` function is not supported by this implementation.
- EPERM** The *cmd* argument is `SETACL` and the effective user ID of the caller does not match the owner of the file.
- EROFS** The *cmd* argument is `SETACL` and the file specified by *pathp* resides on a file system that is mounted read-only.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

SEE ALSO

`getfacl(1)` , `setfacl(1)` , `aclcheck(3)` , `aclsort(3)`

NAME	adjtime – correct the time to allow synchronization of the system clock
SYNOPSIS	<pre>#include <sys/time.h> int adjtime(struct timeval *delta, struct timeval *olddelta);</pre>
DESCRIPTION	<p>The adjtime() function adjusts the system's notion of the current time as returned by gettimeofday(3C), advancing or retarding it by the amount of time specified in the <code>struct timeval</code> pointed to by <i>delta</i>.</p> <p>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. 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.</p> <p>If <i>delta</i> is 0, then <i>olddelta</i> returns the status of the effects of the previous adjtime() call with no effect on the time correction as a result of this call. If <i>olddelta</i> is not a null pointer, then the structure it points to will contain, upon successful return, the number of seconds and/or microseconds still to be corrected from the earlier call. If <i>olddelta</i> is a null pointer, the corresponding information will not be returned.</p> <p>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.</p> <p>Only the super-user may adjust the time of day.</p> <p>The adjustment value will be silently rounded to the resolution of the system clock.</p>
RETURN VALUES	Upon successful completion, adjtime() returns 0. Otherwise, it returns -1 and sets <code>errno</code> to indicate the error.
ERRORS	<p>The adjtime() function will fail if:</p> <p>EFAULT The <i>delta</i> or <i>olddelta</i> argument points outside the process's allocated address space, or <i>olddelta</i> points to a region of the process's allocated address space that is not writable.</p> <p>EINVAL The <code>tv_usec</code> member of <i>delta</i> is not within valid range (-1000000 to 1000000).</p> <p>EPERM The effective user of the calling process is not super-user. Additionally, the adjtime() function will fail for 32-bit interfaces if:</p>

EOverflow The size of the `tv_sec` member of the `timeval` structure pointed to by *olddelta* is too small to contain the correct number of seconds.

SEE ALSO `date(1)`, `gettimeofday(3C)`

NAME	alarm – set a process alarm clock
SYNOPSIS	<pre>#include <unistd.h> unsigned int alarm(unsigned int sec);</pre>
DESCRIPTION	<p>The alarm() function instructs the alarm clock of the calling process to send the signal SIGALRM to the calling process after the number of real time seconds specified by <i>sec</i> have elapsed (see signal(3C)).</p> <p>Alarm requests are not stacked; successive calls reset the alarm clock of the calling process.</p> <p>If <i>sec</i> is 0, any previously made alarm request is canceled.</p> <p>The fork(2) function sets the alarm clock of a new process to 0. A process created by the exec family of routines inherits the time left on the old process's alarm clock.</p> <p>Calling alarm() in a multithreaded process linked with -lthread (Solaris threads) and not with -lpthread (POSIX threads) currently behaves in the following fashion:</p> <ul style="list-style-type: none"> ■ if the calling thread is a bound thread, the resulting SIGALRM is delivered to the bound thread's LWP, i.e. to the calling thread. There is a bug currently that this signal is not maskable via thr_sigsetmask(3T) on this bound thread. ■ if the calling thread is an unbound thread, the resulting SIGALRM is sent to the LWP on which the thread was running when it issued the call to alarm(). This is neither a per-process semantic, nor a per-thread semantic, since the LWP could change threads after the call to alarm() but before the SIGALRM delivery, causing some other thread to get it possibly. Hence this is basically a bug. <p>The above documents current behavior and the bugs are not going to be fixed since the above semantics are going to be discontinued in the next release.</p> <p>The semantic for Solaris threads will move to the per-process semantic specified by POSIX (see standards(5)) at this future date. New applications should not rely on the per-thread semantic of alarm(), since this semantic will become obsolete.</p> <p>In a process linked with -lpthread (whether or not it is also linked with -lthread), the semantics of alarm() are per-process; the resulting SIGALRM is sent to the process, and not necessarily to the calling thread. This semantic will be supported in the future.</p>

This semantic is obtainable by simply linking with `-lpthread`. One can continue to use Solaris thread interfaces by linking with both `-lpthread` and `-lthread`.

RETURN VALUES

The **alarm()** function returns the amount of time previously remaining in the alarm clock of the calling process.

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

exec(2), **fork(2)**, **pause(2)**, **signal(3C)**, **thr_sigsetmask(3T)**, **attributes(5)**, **standards(5)**

NAME	audit – write a record to the audit log
SYNOPSIS	<pre>cc [<i>flag ...</i>] <i>file ...</i> -lbsm -lsocket -lnsl -lintl [<i>library...</i>] #include <sys/param.h> #include <bsm/audit.h> int audit(caddr_t <i>record</i>, int <i>length</i>);</pre>
DESCRIPTION	<p>The audit() function is used to write a record to the system audit log. The data pointed to by <i>record</i> is written to the log after a minimal consistency check, with the <i>length</i> parameter specifying the size of the record in bytes. The data should be a well-formed audit record as described by audit.log(4).</p> <p>The kernel validates the record header token type and length, and sets the time stamp value before writing the record to the audit log. The kernel does not do any preselection for user-level generated events. If the audit policy is set to include sequence or trailer tokens, the kernel will append them to the record.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and errno is set to indicate the error.
ERRORS	<p>The audit() function will fail if:</p> <p>EFAULT The <i>record</i> argument points outside the process's allocated address space.</p> <p>EINVAL The record header token ID is invalid or the length is either less than the header token size or greater than MAXAUDITDATA.</p> <p>EPERM The process's effective user ID is not super-user.</p>
USAGE	Only the super-user may successfully execute this call.
SEE ALSO	bsmconv(1M) , auditd(1M) , auditon(2) , auditsvc(2) , getaudit(2) , audit.log(4)
NOTES	The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See bsmconv(1M) for more information.

NAME	auditon - manipulate auditing
SYNOPSIS	<pre>cc [<i>flag ...</i>] <i>file ...</i> -lbsm -lsocket -lnsl -lintl [<i>library ...</i>] #include <sys/param.h> #include <bsm/audit.h> int auditon(int <i>cmd</i>, caddr_t <i>data</i>, int <i>length</i>);</pre>
DESCRIPTION	<p>The auditon() function performs various audit subsystem control operations. The <i>cmd</i> argument designates the particular audit control command. The <i>data</i> argument is a pointer to command-specific data. The <i>length</i> argument is the length in bytes of the command-specific data.</p> <p>The following commands are supported:</p> <p>A_GETCOND Return the system audit on/off/disabled condition in the integer long pointed to by <i>data</i>. The following values may be returned:</p> <ul style="list-style-type: none"> AUC_AUDITING Auditing has been turned on. AUC_NOAUDIT Auditing has been turned off. AUC_DISABLED Auditing package installed, not turned on. <p>A_SETCOND Set the system's audit on/off condition to the value in the integer long pointed to by <i>data</i>. The BSM audit module must be enabled by bsmconv(1M) before auditing can be turned on. The following audit states may be set:</p> <ul style="list-style-type: none"> AUC_AUDITING Turns on audit record generation. AUC_NOAUDIT Turns off audit record generation. <p>A_GETCLASS Return the event to class mapping for the designated audit event. The <i>data</i> argument points to the <code>au_evclass_map</code> structure containing the event number. The preselection class mask is returned in the same structure.</p> <p>A_SETCLASS Set the event class preselection mask for the designated audit event. The <i>data</i> argument points to the <code>au_evclass_map</code> structure containing the event number and class mask.</p> <p>A_GETKMASK Return the kernel preselection mask in the <code>au_mask</code> structure pointed to by <i>data</i>. This is the mask used to preselect non-attributable audit events.</p>

A_SETKMASK	Set the kernel preselection mask. The <i>data</i> argument points to the <code>au_mask</code> structure containing the class mask. This is the mask used to preselect non-attributable audit events.
A_GETPINFO	Return the audit ID, preselection mask, terminal ID and audit session ID of the specified process in the <code>auditpinfo</code> structure pointed to by <i>data</i> .
A_SETPMASK	Set the preselection mask of the specified process. The <i>data</i> argument points to the <code>auditpinfo</code> structure containing the process ID and the preselection mask. The other fields of the structure are ignored and should be set to <code>NULL</code> .
A_SETUMASK	Set the preselection mask for all processes with the specified audit ID. The <i>data</i> argument points to the <code>auditinfo</code> structure containing the audit ID and the preselection mask. The other fields of the structure are ignored and should be set to <code>NULL</code> .
A_SETSMASK	Set the preselection mask for all processes with the specified audit session ID. The <i>data</i> argument points to the <code>auditinfo</code> structure containing the audit session ID and the preselection mask. The other fields of the structure are ignored and should be set to <code>NULL</code> .
A_GETQCTRL	Return the kernel audit queue control parameters. These control the high and low water marks of the number of audit records allowed in the audit queue. The high water mark is the maximum allowed number of undelivered audit records. The low water mark determines when threads blocked on the queue are wakened. Another parameter controls the size of the data buffer used by <code>auditsvc(2)</code> to write data to the audit trail. There is also a parameter that specifies a maximum delay before data is attempted to be written to the audit trail. The audit queue parameters are returned in the <code>au_qctrl</code> structure pointed to by <i>data</i> .
A_SETQCTRL	Set the kernel audit queue control parameters as described above in the <code>A_GETQCTRL</code> command. The <i>data</i> argument points to the <code>au_qctrl</code> structure containing the audit queue control parameters. The default and maximum values 'A/B' for the audit queue control parameters are:

high water	100/10000 (audit records)
-------------------	---------------------------

	low water	10/1024 (audit records)
	output buffer size	1024/1048576 (bytes)
	delay	20/20000 (hundredths second)
A_GETCWD	Return the current working directory as kept by the audit subsystem. This is a path anchored on the real root, rather than on the active root. The <i>data</i> argument points to a buffer into which the path is copied. The <i>length</i> argument is the length of the buffer.	
A_GETCAR	Return the current active root as kept by the audit subsystem. This path may be used to anchor an absolute path for a path token generated by an application. The <i>data</i> argument points to a buffer into which the path is copied. The <i>length</i> argument is the length of the buffer.	
A_GETSTAT	Return the system audit statistics in the <code>audit_stat</code> structure pointed to by <i>data</i> .	
A_SETSTAT	Reset system audit statistics values. The kernel statistics value is reset if the corresponding field in the statistics structure pointed to by the <i>data</i> argument is <code>CLEAR_VAL</code> . Otherwise, the value is not changed.	
A_SETFSIZE	Set the maximum size of an audit trail file. When the audit file reaches the designated size, it is closed and a new file started. If the maximum size is unset, the audit trail file generated by <code>auditsvc()</code> will grow to the size of the file system. The <i>data</i> argument points to the <code>au_fstat_t</code> structure containing the maximum audit file size in bytes. The size can not be set less than <code>0x80000</code> bytes.	
A_GETFSIZE	Return the maximum audit file size and current file size in the <code>au_fstat_t</code> structure pointed to by the <i>data</i> argument.	
A_GETPOLICY	Return the audit policy flags in the integer long pointed to by <i>data</i> .	

A_SETPOLICY	Set the audit policy flags to the values in the integer long pointed to by <i>data</i> . The following policy flags are recognized:
AUDIT_CNT	Do not suspend processes when audit storage is full or inaccessible. The default action is to suspend processes until storage becomes available.
AUDIT_AHLT	Halt the machine when a non-attributable audit record can not be delivered. The default action is to count the number of events that could not be recorded.
AUDIT_ARGV	Include in the audit record the argument list for a member of the <i>exec</i> family of functions (see <i>exec(2)</i>). The default action is not to include this information.
AUDIT_ARGE	Include the environment variables for the <i>execv(2)</i> function in the audit record. The default action is not to include this information.
AUDIT_SEQ	Add a <i>sequence</i> token to each audit record. The default action is not to include it.
AUDIT_TRAIL	Append a <i>trailer</i> token to each audit record. The default action is not to include it.
AUDIT_GROUP	Include the supplementary groups list in audit records. The default action is not to include it.
AUDIT_PATH	Include secondary paths in audit records. Examples of secondary paths are dynamically loaded shared library modules and the command shell path for executable scripts. The default action is to include only the primary path from the system call.

RETURN VALUES

Upon successful completion, **auditon()** returns 0. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The **auditon()** function will fail if:

EFAULT The copy of data to/from the kernel failed.

EINVAL One of the arguments was illegal, or BSM has not been installed.

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

USAGE The **auditon()** function may be invoked only by processes with super-user privileges.

SEE ALSO **auditconfig(1M)**, **auditd(1M)**, **bsmconv(1M)**, **audit(2)**, **auditsvc(2)**, **exec(2)**, **audit.log(4)**

NOTES The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See **bsmconv(1M)** for more information.

NAME	auditsvc – write audit log to specified file descriptor										
SYNOPSIS	<pre>cc [<i>flag</i> ...] <i>file</i>... -lbsm -lsocket -lnsl -lintl [<i>library</i> ...] #include <sys/param.h> #include <bsm/audit.h> int auditsvc(int <i>fd</i>, int <i>limit</i>);</pre>										
DESCRIPTION	<p>The auditsvc() function 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 <i>fd</i> argument is a file descriptor that identifies the audit file. Applications should open this file for writing before calling auditsvc().</p> <p>The <i>limit</i> argument specifies the number of free blocks that must be available in the audit file system, and causes auditsvc() to return when the 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.</p> <p>The auditsvc() function does not return until one of the following conditions occurs:</p> <ul style="list-style-type: none"> ■ 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 <i>limit</i>), has been reached. 										
RETURN VALUES	The auditsvc() function returns only on an error.										
ERRORS	<p>The auditsvc() function will fail if:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 10px;">EAGAIN</td> <td>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.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 10px;">EBADF</td> <td>The <i>fd</i> argument is not a valid descriptor open for writing.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 10px;">EBUSY</td> <td>A second process attempted to perform this call.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 10px;">EFBIG</td> <td>An attempt was made to write a file that exceeds the process's file size limit or the maximum file size.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 10px;">EINTR</td> <td>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 sigvec(3B)). The signal(3C) function sets this bit for any signal it catches.</td> </tr> </table>	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	The <i>fd</i> argument is not a valid descriptor open for writing.	EBUSY	A second process attempted to perform this call.	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 sigvec(3B)). The signal(3C) function sets this bit for any signal it catches.
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	The <i>fd</i> argument is not a valid descriptor open for writing.										
EBUSY	A second process attempted to perform this call.										
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 sigvec(3B)). The signal(3C) function sets this bit for any signal it catches.										

EINVAL	Auditing is disabled (see <code>auditon(2)</code>), or the <i>fd</i> argument 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	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; or 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 user ID is not super-user.
EWOULDBLOCK	The file was marked for 4.2 BSD-style non-blocking I/O, and no data could be written immediately.
USAGE	Only processes with an effective user ID of super-user may execute this call successfully.
SEE ALSO	<code>auditd(1M)</code> , <code>bsmconv(1M)</code> , <code>audit(2)</code> , <code>auditon(2)</code> , <code>sigvec(3B)</code> , <code>audit.log</code> (4)
NOTES	The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See <code>bsmconv(1M)</code> for more information.

NAME	brk, sbrk – change the amount of space allocated for the calling process's data segment
SYNOPSIS	<pre>#include <unistd.h> int brk(void * <i>ends</i>); void * sbrk(intptr_t <i>incr</i>);</pre>
DESCRIPTION	<p>The brk() and sbrk() functions are used to change dynamically the amount of space allocated for the calling process's data segment (see exec(2)). The change is made by resetting the process's break value and allocating the appropriate amount of space. The break value is the address of the first location beyond the end of the data segment. The amount of allocated space increases as the break value increases. Newly allocated space is set to zero. If, however, the same memory space is reallocated to the same process its contents are undefined.</p> <p>When a program begins execution using execve() the break is set at the highest location defined by the program and data storage areas.</p> <p>The getrlimit(2) function may be used to determine the maximum permissible size of the <i>data</i> segment; it is not possible to set the break beyond the <code>rlim_max</code> value returned from a call to getrlimit(), that is to say, “<code>end + rlim.rlim_max</code>.” See end(3C).</p> <p>The brk() function sets the break value to <i>ends</i> and changes the allocated space accordingly.</p> <p>The sbrk() function adds <i>incr</i> function bytes to the break value and changes the allocated space accordingly. The <i>incr</i> function can be negative, in which case the amount of allocated space is decreased.</p>
RETURN VALUES	<p>Upon successful completion, brk() returns 0. Otherwise, it returns -1 and sets <code>errno</code> to indicate the error.</p> <p>Upon successful completion, sbrk() returns the prior break value. Otherwise, it returns <code>(void *)-1</code> and sets <code>errno</code> to indicate the error.</p>
ERRORS	<p>The brk() and sbrk() functions will fail and no additional memory will be allocated if:</p> <p>ENOMEM The data segment size limit as set by setrlimit() (see getrlimit(2)) would be exceeded; the maximum possible size of a data segment (compiled into the system) would be exceeded; insufficient space exists in the swap area to support the expansion; or the new break value would extend into an area of the address space defined by some previously established mapping (see mmap(2)).</p>

EAGAIN Total amount of system memory available for private pages is temporarily insufficient. This may occur even though the space requested was less than the maximum data segment size (see `ulimit(2)`).

USAGE The behavior of `brk()` and `sbrk()` is unspecified if an application also uses any other memory functions (such as `malloc(3C)`, `mmap(2)`, `free(3C)`). The `brk()` and `sbrk()` functions have been used in specialized cases where no other memory allocation function provided the same capability. The use of `mmap(2)` is now preferred because it can be used portably with all other memory allocation functions and with any function that uses other allocation functions.

It is unspecified whether the pointer returned by `sbrk()` is aligned suitably for any purpose.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `exec(2)`, `getrlimit(2)`, `mmap(2)`, `shmop(2)`, `ulimit(2)`, `end(3C)`, `free(3C)`, `malloc(3C)`

NOTES The value of `incr` may be adjusted by the system before setting the new break value. Upon successful completion, the implementation guarantees a minimum of `incr` bytes will be added to the data segment if `incr` is a positive value. If `incr` is a negative value, a maximum of `incr` bytes will be removed from the data segment. This adjustment may not be necessary for all machine architectures.

The value of the arguments to both `brk()` and `sbrk()` are rounded up for alignment with eight-byte boundaries.

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, fchdir – change working directory
SYNOPSIS	<pre>#include <unistd.h> int chdir(const char * path); int fchdir(int fildes);</pre>
DESCRIPTION	<p>The chdir() and fchdir() functions cause a directory pointed to by <i>path</i> or <i>fildes</i> to become the current working directory. The starting point for path searches for path names not beginning with / (slash). The <i>path</i> argument points to the path name of a directory. The <i>fildes</i> argument is an open file descriptor of a directory.</p> <p>For a directory to become the current directory, a process must have execute (search) access to the directory.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, the current working directory is unchanged, and <code>errno</code> is set to indicate the error.
ERRORS	<p>The chdir() function will fail if:</p> <p>EACCES Search permission is denied for any component of the path name.</p> <p>EFAULT The <i>path</i> argument points to an illegal address.</p> <p>EINTR A signal was caught during the execution of the chdir() function.</p> <p>EIO An I/O error occurred while reading from or writing to the file system.</p> <p>ELOOP Too many symbolic links were encountered in translating <i>path</i> .</p> <p>ENAMETOOLONG The length of the <i>path</i> argument exceeds <code>PATH_MAX</code> , or the length of a <i>path</i> component exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.</p> <p>ENOENT Either a component of the path prefix or the directory named by <i>path</i> does not exist or is a null pathname.</p> <p>ENOLINK The <i>path</i> argument points to a remote machine and the link to that machine is no longer active.</p>

ENOTDIR	A component of the path name is not a directory.
The fchdir() function will fail if:	
EACCES	Search permission is denied for <i>filde</i> .
EBADF	The <i>filde</i> argument is not an open file descriptor.
EINTR	A signal was caught during the execution of the fchdir() function.
EIO	An I/O error occurred while reading from or writing to the file system.
ENOLINK	The <i>filde</i> argument points to a remote machine and the link to that machine is no longer active.
ENOTDIR	The open file descriptor <i>filde</i> does not refer to a directory.

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	chdir() is Async-Signal-Safe

SEE ALSO

chroot(2) , **attributes(5)**

NAME	chmod, fchmod – change access permission mode of file																																													
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/stat.h> int chmod(const char * <i>path</i>, mode_t <i>mode</i>); int fchmod(int <i>fd</i>, mode_t <i>mode</i>);</pre>																																													
DESCRIPTION	<p>The chmod() and fchmod() functions set the access permission portion of the mode of the file whose name is given by <i>path</i> or referenced by the open file descriptor <i>fd</i> to the bit pattern contained in <i>mode</i>. Access permission bits are interpreted as follows:</p> <table border="0"> <tr> <td style="padding-right: 20px;">S_ISUID</td> <td style="padding-right: 20px;">04000</td> <td>Set user ID on execution.</td> </tr> <tr> <td>S_ISGID</td> <td>020#0</td> <td>Set group ID on execution if # is 7, 5, 3, or 1. Enable mandatory file/record locking if # is 6, 4, 2, or 0.</td> </tr> <tr> <td>S_ISVTX</td> <td>01000</td> <td>Save text image after execution.</td> </tr> <tr> <td>S_IRWXU</td> <td>00700</td> <td>Read, write, execute by owner.</td> </tr> <tr> <td>S_IRUSR</td> <td>00400</td> <td>Read by owner.</td> </tr> <tr> <td>S_IWUSR</td> <td>00200</td> <td>Write by owner.</td> </tr> <tr> <td>S_IXUSR</td> <td>00100</td> <td>Execute (search if a directory) by owner.</td> </tr> <tr> <td>S_IRWXG</td> <td>00070</td> <td>Read, write, execute by group.</td> </tr> <tr> <td>S_IRGRP</td> <td>00040</td> <td>Read by group.</td> </tr> <tr> <td>S_IWGRP</td> <td>00020</td> <td>Write by group.</td> </tr> <tr> <td>S_IXGRP</td> <td>00010</td> <td>Execute by group.</td> </tr> <tr> <td>S_IRWXO</td> <td>00007</td> <td>Read, write, execute (search) by others.</td> </tr> <tr> <td>S_IROTH</td> <td>00004</td> <td>Read by others.</td> </tr> <tr> <td>S_IWOTH</td> <td>00002</td> <td>Write by others.</td> </tr> <tr> <td>S_IXOTH</td> <td>00001</td> <td>Execute by others.</td> </tr> </table> <p>Modes are constructed by the bitwise OR operation of the access permission bits.</p> <p>The effective user ID of the process must match the owner of the file or the process must have the appropriate privilege to change the mode of a file.</p> <p>If the process is not a privileged process and the file is not a directory, mode bit 01000 (save text image on execution) is cleared.</p>	S_ISUID	04000	Set user ID on execution.	S_ISGID	020#0	Set group ID on execution if # is 7, 5, 3, or 1. Enable mandatory file/record locking if # is 6, 4, 2, or 0.	S_ISVTX	01000	Save text image after execution.	S_IRWXU	00700	Read, write, execute by owner.	S_IRUSR	00400	Read by owner.	S_IWUSR	00200	Write by owner.	S_IXUSR	00100	Execute (search if a directory) by owner.	S_IRWXG	00070	Read, write, execute by group.	S_IRGRP	00040	Read by group.	S_IWGRP	00020	Write by group.	S_IXGRP	00010	Execute by group.	S_IRWXO	00007	Read, write, execute (search) by others.	S_IROTH	00004	Read by others.	S_IWOTH	00002	Write by others.	S_IXOTH	00001	Execute by others.
S_ISUID	04000	Set user ID on execution.																																												
S_ISGID	020#0	Set group ID on execution if # is 7, 5, 3, or 1. Enable mandatory file/record locking if # is 6, 4, 2, or 0.																																												
S_ISVTX	01000	Save text image after execution.																																												
S_IRWXU	00700	Read, write, execute by owner.																																												
S_IRUSR	00400	Read by owner.																																												
S_IWUSR	00200	Write by owner.																																												
S_IXUSR	00100	Execute (search if a directory) by owner.																																												
S_IRWXG	00070	Read, write, execute by group.																																												
S_IRGRP	00040	Read by group.																																												
S_IWGRP	00020	Write by group.																																												
S_IXGRP	00010	Execute by group.																																												
S_IRWXO	00007	Read, write, execute (search) by others.																																												
S_IROTH	00004	Read by others.																																												
S_IWOTH	00002	Write by others.																																												
S_IXOTH	00001	Execute by others.																																												

If neither the process is privileged, nor the file's group is a member of the process's supplementary group list, and the effective group ID of the process does not match the group ID of the file, mode bit 02000 (set group ID on execution) is cleared.

If a directory is writable and has `S_ISVTX` (the sticky bit) set, files within that directory can be removed or renamed only if one or more of the following is true (see `unlink(2)` and `rename(2)`):

- the user owns the file
- the user owns the directory
- the file is writable by the user
- the user is a privileged user

If a directory has the set group ID bit set, a given file created within that directory will have the same group ID as the directory, if that group ID is part of the group ID set of the process that created the file. Otherwise, the newly created file's group ID will be set to the effective group ID of the creating process.

If the mode bit 02000 (set group ID on execution) is set and the mode bit 00010 (execute or search by group) is not set, mandatory file/record locking will exist on a regular file. This may affect future calls to `open(2)`, `creat(2)`, `read(2)`, and `write(2)` on this file.

Upon successful completion, `chmod()` and `fchmod()` mark for update the `st_ctime` field of the file.

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned, the file mode is unchanged, and `errno` is set to indicate the error.

ERRORS

The `chmod()` function will fail if:

EACCES	Search permission is denied on a component of the path prefix of <i>path</i> .
EFAULT	The <i>path</i> argument points to an illegal address.
EINTR	A signal was caught during execution of the function.
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 <i>path</i> argument exceeds <code>PATH_MAX</code> , or the length of a <i>path</i> component exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.
ENOENT	Either a component of the path prefix or the file referred to by <i>path</i> does not exist or is a null pathname.
ENOLINK	The <i>fildev</i> argument points to a remote machine and the link to that machine is no longer active.
ENOTDIR	A component of the prefix of <i>path</i> is not a directory.
EPERM	The effective user ID does not match the owner of the file and is not super-user.
EROFS	The file referred to by <i>path</i> resides on a read-only file system.
The <code>fchmod()</code> function will fail if:	
EBADF	The <i>fildev</i> argument is not an open file descriptor
EIO	An I/O error occurred while reading from or writing to the file system.
EINTR	A signal was caught during execution of the <code>fchmod()</code> function.
ENOLINK	The <i>path</i> argument points to a remote machine and the link to that machine is no longer active.
EPERM	The effective user ID does not match the owner of the file and the process does not have appropriate privilege.
EROFS	The file referred to by <i>fildev</i> resides on a read-only file system.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	<code>chmod()</code> is Async-Signal-Safe

SEE ALSO `chmod(1)`, `chown(2)`, `creat(2)`, `fcntl(2)`, `mknod(2)`, `open(2)`, `read(2)`, `rename(2)`, `stat(2)`, `write(2)`, `mkfifo(3C)`, `attributes(5)`, `stat(5)`

System Interface Guide

NOTES If you use `chmod()` to change the file group owner permissions on a file with ACL entries, both the file group owner permissions and the ACL mask are changed to the new permissions. Be aware that the new ACL mask permissions may change the effective permissions for additional users and groups who have ACL entries on the file.

NAME	chown, lchown, fchown – change owner and group of a file
SYNOPSIS	<pre>#include <unistd.h> #include <sys/types.h> int chown(const char * path, uid_t owner, gid_t group); int lchown(const char * path, uid_t owner, gid_t group); int fchown(int fildes, uid_t owner, gid_t group);</pre>
DESCRIPTION	<p>The chown() function sets the owner ID and group ID of the file specified by <i>path</i> or referenced by the open file descriptor <i>fildes</i> to <i>owner</i> and <i>group</i> respectively. If <i>owner</i> or <i>group</i> is specified as <i>-1</i>, chown() does not change the corresponding ID of the file.</p> <p>The lchown() function sets the owner ID and group ID of the named file in the same manner as chown(), unless the named file is a symbolic link. In this case, lchown() changes the ownership of the symbolic link file itself, while chown() changes the ownership of the file or directory to which the symbolic link refers.</p> <p>If chown(), lchown(), or fchown() is invoked by a process other than super-user, the set-user-ID and set-group-ID bits of the file mode, <i>S_ISUID</i> and <i>S_ISGID</i> respectively, are cleared (see chmod(2)).</p> <p>The operating system provides a configuration option, { <i>_POSIX_CHOWN_RESTRICTED</i> }, to restrict ownership changes for the chown(), lchown(), and fchown() functions. When { <i>_POSIX_CHOWN_RESTRICTED</i> } is not in effect, either the effective user ID of the process must match the owner of the file or the process must be the super-user to change the ownership of a file. When { <i>_POSIX_CHOWN_RESTRICTED</i> } is in effect (the default behavior), the chown(), lchown(), and fchown() functions, for users other than super-user, prevent the owner of the file from changing the owner ID of the file and restrict the change of the group of the file to the list of supplementary group IDs. To set this configuration option, include the following line in <i>/etc/system</i>:</p> <pre>set rstchown = 1</pre> <p>To disable this option, include the following line in <i>/etc/system</i>:</p> <pre>set rstchown = 0</pre> <p>See system(4) and fpathconf(2).</p> <p>Upon successful completion, chown(), fchown() and lchown() mark for update the <i>st_ctime</i> field of the file.</p>

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned, the owner and group of the named file remain unchanged, and `errno` is set to indicate the error.

ERRORS

The `chown()` and `lchown()` functions will fail if:

EACCES	Search permission is denied on a component of the path prefix of <i>path</i> .
EFAULT	The <i>path</i> argument points to an illegal address.
EINTR	A signal was caught during the execution of the <code>chown()</code> or <code>lchown()</code> function.
EINVAL	The <i>group</i> or <i>owner</i> argument is out of range.
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 <i>path</i> argument exceeds { <code>PATH_MAX</code> }, or the length of a <i>path</i> component exceeds { <code>NAME_MAX</code> } while { <code>_POSIX_NO_TRUNC</code> } is in effect.
ENOLINK	The <i>path</i> argument points to a remote machine and the link to that machine is no longer active.
ENOENT	Either a component of the path prefix or the file referred to by <i>path</i> does not exist or is a null pathname.
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 or the process is not the super-user and <code>_POSIX_CHOWN_RESTRICTED</code> indicates that such privilege is required.
EROFS	The named file resides on a read-only file system.
The <code>fchown()</code> function will fail if:	
EBADF	The <i>fdes</i> argument is not an open file descriptor.
EIO	An I/O error occurred while reading from or writing to the file system.

EINTR	A signal was caught during execution of the function.
ENOLINK	The <i>fildev</i> argument points to a remote machine and the link to that machine is no longer active.
EINVAL	The <i>group</i> or <i>owner</i> argument is out of range.
EPERM	The effective user ID does not match the owner of the file, or the process is not the super-user and <code>_POSIX_CHOWN_RESTRICTED</code> indicates that such privilege is required.
EROFS	The named file referred to by <i>fildev</i> resides on a read-only file system.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	chown() is Async-Signal-Safe

SEE ALSO

`chgrp(1)`, `chown(1)`, `chmod(2)`, `fpathconf(2)`, `system(4)`, `attributes(5)`

NAME	chroot, fchroot – change root directory
SYNOPSIS	<pre>#include <unistd.h> int chroot(const char * path); int fchroot(int fildes);</pre>
DESCRIPTION	<p>The chroot() and fchroot() functions cause a directory to become the root directory, the starting point for path searches for path names beginning with / (slash). The user's working directory is unaffected by the chroot() and fchroot() functions.</p> <p>The <i>path</i> argument points to a path name naming a directory. The <i>fildes</i> argument to fchroot() is the open file descriptor of the directory which is to become the root.</p> <p>The effective user ID of the process must be super-user to change the root directory. While it is always possible to change to the system root using the fchroot() function, it is not guaranteed to succeed in any other case, even should <i>fildes</i> be valid in all respects.</p> <p>The “.” entry in the root directory is interpreted to mean the root directory itself. Therefore, “.” cannot be used to access files outside the subtree rooted at the root directory. Instead, fchroot() can be used to reset the root to a directory that was opened before the root directory was changed.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, the root directory remains unchanged, and <code>errno</code> is set to indicate the error.
ERRORS	<p>The chroot() function will fail if:</p> <p>EACCES Search permission is denied for a component of the path prefix of <i>dirname</i>, or search permission is denied for the directory referred to by <i>dirname</i>.</p> <p>EBADF The descriptor is not valid.</p> <p>EFAULT The <i>path</i> argument points to an illegal address.</p> <p>EINVAL The fchroot() function attempted to change to a directory the is not the system root and external circumstances do not allow this.</p> <p>EINTR A signal was caught during the execution of the chroot() function.</p> <p>EIO An I/O error occurred while reading from or writing to the file system.</p>

ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
ENAMETOOLONG	The length of the <i>path</i> argument exceeds <code>PATH_MAX</code> , or the length of a <i>path</i> component exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.
ENOENT	The named directory does not exist or is a null pathname.
ENOLINK	The <i>path</i> argument points to a remote machine and the link to that machine is no longer active.
ENOTDIR	Any component of the path name is not a directory.
EPERM	The effective user of the calling process is not super-user.

SEE ALSO `chroot(1M)` , `chdir(2)`

WARNINGS The only use of `fchroot()` that is appropriate is to change back to the system root.

NAME	close – close a file descriptor
SYNOPSIS	<pre>#include <unistd.h> int close(int <i>fdes</i>);</pre>
DESCRIPTION	<p>The close() function will deallocate the file descriptor indicated by <i>fdes</i>. To deallocate means to make the file descriptor available for return by subsequent calls to open(2) or other functions that allocate file descriptors. All outstanding record locks owned by the process on the file associated with the file descriptor will be removed (that is, unlocked).</p> <p>If close() is interrupted by a signal that is to be caught, it will return <code>-1</code> with <code>errno</code> set to <code>EINTR</code> and the state of <i>fdes</i> is unspecified.</p> <p>When all file descriptors associated with a pipe or FIFO special file are closed, any data remaining in the pipe or FIFO will be discarded.</p> <p>When all file descriptors associated with an open file description have been closed the open file description will be freed.</p> <p>If the link count of the file is 0, when all file descriptors associated with the file are closed, the space occupied by the file will be freed and the file will no longer be accessible.</p> <p>If a STREAMS-based (see intro(2)) <i>fdes</i> is closed and the calling process was previously registered to receive a <code>SIGPOLL</code> signal (see signal(3C)) for events associated with that STREAM (see <code>I_SETSIG</code> in streamio(7I)), the calling process will be unregistered for events associated with the STREAM. The last close() for a STREAM causes the STREAM associated with <i>fdes</i> to be dismantled. If <code>O_NONBLOCK</code> and <code>O_NDELAY</code> are not set and there have been no signals posted for the STREAM, and if there is data on the module's write queue, close() waits up to 15 seconds (for each module and driver) for any output to drain before dismantling the STREAM. The time delay can be changed via an <code>I_SETCLTIME</code> ioctl(2) request (see streamio(7I)). If the <code>O_NONBLOCK</code> or <code>O_NDELAY</code> flag is set, or if there are any pending signals, close() does not wait for output to drain, and dismantles the STREAM immediately.</p> <p>If <i>fdes</i> is associated with one end of a pipe, the last close() causes a hangup to occur on the other end of the pipe. In addition, if the other end of the pipe has been named by fattach(3C), then the last close() forces the named end to be detached by fdetach(3C). If the named end has no open file descriptors associated with it and gets detached, the STREAM associated with that end is also dismantled.</p> <p>If <i>fdes</i> refers to the master side of a pseudo-terminal, a <code>SIGHUP</code> signal is sent to the process group, if any, for which the slave side of the pseudo-terminal is</p>

the controlling terminal. It is unspecified whether closing the master side of the pseudo-terminal flushes all queued input and output.

If *fdes* refers to the slave side of a STREAMS-based pseudo-terminal, a zero-length message may be sent to the master.

If *fdes* refers to a socket, **close()** causes the socket to be destroyed. If the socket is connection-mode, and the `SOCK_LINGER` option is set for the socket, and the socket has untransmitted data, then **close()** will block for up to the current linger interval until all data is transmitted.

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned and `errno` is set to indicate the error.

ERRORS

The **close()** function will fail if:

- EBADF** The *fdes* argument is not a valid file descriptor.
- EINTR** The **close()** function was interrupted by a signal.
- ENOLINK** The *fdes* argument is on a remote machine and the link to that machine is no longer active.
- ENOSPC** There was no free space remaining on the device containing the file.

The **close()** function may fail if:

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

USAGE

An application that used the `stdio` function **fopen(3S)** to open a file should use the corresponding **fclose(3S)** function rather than **close()**.

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

intro(2), **creat(2)**, **dup(2)**, **exec(2)**, **fcntl(2)**, **ioctl(2)**, **open(2)**, **pipe(2)**, **fattach(3C)**, **fclose(3S)**, **fdetach(3C)**, **fopen(3S)**, **signal(3C)**, **attributes(5)**, **signal(5)**, **streamio(7I)**

NAME	creat – create a new file or rewrite an existing one
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/stat.h> #include <fcntl.h> int creat(const char *path, mode_t mode);</pre>
DESCRIPTION	<p>The creat() function creates a new ordinary file or prepares to rewrite an existing file named by the path name pointed to by <i>path</i>.</p> <p>If the file exists, the length is truncated to 0 and the mode and owner are unchanged.</p> <p>If the file does not exist the file's owner ID is set to the effective user ID of the process. The group ID of the file is set to the effective group ID of the process, or if the <code>S_ISGID</code> bit is set in the parent directory then the group ID of the file is inherited from the parent directory. The access permission bits of the file mode are set to the value of <i>mode</i> modified as follows:</p> <ul style="list-style-type: none"> ■ If the group ID of the new file does not match the effective group ID or one of the supplementary group IDs, the <code>S_ISGID</code> bit is cleared. ■ All bits set in the process's file mode creation mask (see umask(2)) are correspondingly cleared in the file's permission mask. ■ The "save text image after execution bit" of the mode is cleared (see chmod(2) for the values of mode). <p>Upon successful completion, a write-only file descriptor is returned and the file is open for writing, even if the mode does not permit writing. The file pointer is set to the beginning of the file. The file descriptor is set to remain open across <i>exec</i> functions (see fcntl(2)). A new file may be created with a mode that forbids writing.</p> <p>The call <code>creat(path, mode)</code> is equivalent to:</p> <pre>open(path, O_WRONLY O_CREAT O_TRUNC, mode)</pre>
RETURN VALUES	<p>Upon successful completion, a non-negative integer representing the lowest numbered unused file descriptor is returned. Otherwise, <code>-1</code> is returned, no files are created or modified, and <code>errno</code> is set to indicate the error.</p>
ERRORS	<p>The creat() function will fail:</p> <p>EACCES Search permission is denied on a component of the path prefix; the file does not exist and the directory in which the file is to be created does not permit writing; or the file exists and write permission is denied.</p>

EAGAIN	The file exists, mandatory file/record locking is set, and there are outstanding record locks on the file (see <code>chmod(2)</code>).
EDQUOT	The directory where the new file entry is being placed cannot be extended because the user's quota of disk blocks on that file system has been exhausted, or the user's quota of inodes on the file system where the file is being created has been exhausted.
EFAULT	The <i>path</i> argument points to an illegal address.
EINTR	A signal was caught during the execution of the <code>creat()</code> function.
EISDIR	The named file is an existing directory.
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
EMFILE	The process has too many open files (see <code>getrlimit(2)</code>).
ENFILE	The system file table is full.
ENOENT	A component of the path prefix does not exist, or the path name is null.
ENOLINK	The <i>path</i> argument points to a remote machine and the link to that machine is no longer active.
ENOSPC	The file system is out of inodes.
ENOTDIR	A component of the path prefix is not a directory.
EOVERFLOW	The file is a large file at the time of <code>creat()</code> .
EROFS	The named file resides or would reside on a read-only file system.

USAGE

The `creat()` function has a transitional interface for 64-bit file offsets. See `1f64(5)`.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

chmod(2), close(2), dup(2), fcntl(2), getrlimit(2), lseek(2), open(2),
read(2), umask(2), write(2), attributes(5), largefile(5), lf64(5),
stat(5)

NAME	dup – duplicate an open file descriptor				
SYNOPSIS	<pre>#include <unistd.h> int dup(int <i>fildes</i>);</pre>				
DESCRIPTION	<p>The dup() function returns a new file descriptor having the following in common with the original open file descriptor <i>fildes</i>:</p> <ul style="list-style-type: none"> ■ same open file (or pipe) ■ same file pointer (that is, both file descriptors share one file pointer) ■ same access mode (read, write or read/write). <p>The new file descriptor is set to remain open across <i>exec</i> functions (see fcntl(2)).</p> <p>The file descriptor returned is the lowest one available.</p> <p>The <code>dup(<i>fildes</i>)</code> function call is equivalent to:</p> <pre>fcntl(<i>fildes</i>, F_DUPFD, 0)</pre>				
RETURN VALUES	<p>Upon successful completion, a non-negative integer representing the file descriptor is returned. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.</p>				
ERRORS	<p>The dup() function will fail if:</p> <p>EBADF The <i>fildes</i> argument is not a valid open file descriptor.</p> <p>EINTR A signal was caught during the execution of the dup() function.</p> <p>EMFILE The process has too many open files (see getrlimit(2)).</p> <p>ENOLINK The <i>fildes</i> argument is on a remote machine and the link to that machine is no longer active.</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	<p>close(2), creat(2), exec(2), fcntl(2), getrlimit(2), open(2), pipe(2), dup2(3C), lockf(3C), attributes(5)</p>				

NAME	exec, execl, execv, execl, execve, execlp, execvp – execute a file
SYNOPSIS	<pre>#include <unistd.h> int execl(const char * path, const char * arg0, ..., const char * argn, char * /*NULL*/); int execv(const char * path, char *const argv[]); int execl_e(const char * path, char *const arg0[], ..., const char * argn, char * /*NULL*/, char *const envp[]); int execve(const char * path, char *const argv[], char *const envp[]); int execlp(const char * file, const char * arg0, ..., const char * argn, char * /*NULL*/); int execvp(const char * file, char *const argv[]);</pre>
DESCRIPTION	<p>Each of the functions in the <code>exec</code> family overlays a new process image on an old process. The new process image is constructed from an ordinary, executable file. This file is either an executable object file, or a file of data for an interpreter. There can be no return from a successful call to one of these functions because the calling process image is overlaid by the new process image.</p> <p>An interpreter file begins with a line of the form</p> <pre>#! <i>pathname</i> [<i>arg</i>]</pre> <p>where <i>pathname</i> is the path of the interpreter, and <i>arg</i> is an optional argument. When an interpreter file is executed, the system invokes the specified interpreter. The <i>pathname</i> specified in the interpreter file is passed as <i>arg0</i> to the interpreter. If <i>arg</i> was specified in the interpreter file, it is passed as <i>arg1</i> to the interpreter. The remaining arguments to the interpreter are <i>arg0</i> through <i>argn</i> of the originally <code>exec'd</code> file. The interpreter named by <i>pathname</i> must not be an interpreter file.</p> <p>When a C program is executed, it is called as follows:</p> <pre>int main (int argc, char *argv[], char *envp[]);</pre> <p>where <i>argc</i> is the argument count, <i>argv</i> is an array of character pointers to the arguments themselves, and <i>envp</i> is an array of character pointers to the environment strings. As indicated, <i>argc</i> is at least one, and the first member of the array points to a string containing the name of the file.</p> <p>The arguments <i>arg0</i>, ..., <i>argn</i> point to null-terminated character strings. These strings constitute the argument list available to the new process image. Conventionally at least <i>arg0</i> should be present. It will become the name of the process, as displayed by the <code>ps(1)</code> command. The <i>arg0</i> argument points to a</p>

string that is the same as *path* (or the last component of *path*). The list of argument strings is terminated by a `(char *)0` argument.

The *argv* argument is an array of character pointers to null-terminated strings. These strings constitute the argument list available to the new process image. By convention, *argv* must have at least one member, and it should point to a string that is the same as *path* (or its last component). The *argv* argument is terminated by a null pointer.

The *envp* argument is an array of character pointers to null-terminated strings. These strings constitute the environment for the new process image. The *envp* argument is terminated by a null pointer. For **execl()**, **execv()**, **execvp()**, and **execlp()**, the C run-time start-off routine places a pointer to the environment of the calling process in the global object `extern char **environ`, and it is used to pass the environment of the calling process to the new process.

The *path* argument points to a path name that identifies the new process file.

The *file* argument points to the new process file. If *file* does not contain a slash character, the path prefix for this file is obtained by a search of the directories passed in the PATH environment variable (see **environ(5)**). The environment is supplied typically by the shell. If the new process file is not an executable object file, **execlp()** and **execvp()** use the contents of that file as standard input to the shell. In a standard-conforming application (see **standards(5)**), the *exec* family of functions use `/usr/bin/ksh` (see **ksh(1)**); otherwise, they use `/usr/bin/sh` (see **sh(1)**).

File descriptors open in the calling process remain open in the new process, except for those whose close-on-exec flag is set; (see **fcntl(2)**). For those file descriptors that remain open, the file pointer is unchanged.

Signals that are being caught by the calling process are set to the default disposition in the new process image (see **signal(3C)**). Otherwise, the new process image inherits the signal dispositions of the calling process.

The saved resource limits in the new process image are set to be a copy of the process's corresponding hard and soft resource limits.

If the set-user-ID mode bit of the new process file is set (see **chmod(2)**), the effective user ID of the new process is set to the owner ID of the new process file. Similarly, if the set-group-ID mode bit of the new process file is set, the effective group ID of the new process is set to the group ID of the new process file. The real user ID and real group ID of the new process remain the same as those of the calling process.

If the effective user-ID is `root` or super-user, the set-user-ID and set-group-ID bits will be honored when the process is being controlled by `ptrace`.

The shared memory segments attached to the calling process will not be attached to the new process (see `shmop(2)`). Memory mappings in the calling process are unmapped before the new process begins execution (see `mmap(2)`).

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

Timers created by `timer_create(3R)` are deleted before the new process begins execution.

Any outstanding asynchronous I/O operations may be cancelled.

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

- nice value (see `nice(2)`)
- scheduler class and priority (see `prctl(2)`)
- process ID
- parent process ID
- process group ID
- supplementary group IDs
- `semadj` values (see `semop(2)`)
- session ID (see `exit(2)` and `signal(3C)`)
- trace flag (see `ptrace(2)` request 0)
- time left until an alarm (see `alarm(2)`)
- current working directory
- root directory
- file mode creation mask (see `umask(2)`)
- resource limits (see `getrlimit(2)`)
- `utime`, `stime`, `cutime`, and `cstime` (see `times(2)`)
- file-locks (see `fcntl(2)` and `lockf(3C)`)
- controlling terminal
- process signal mask (see `sigprocmask(2)`)
- pending signals (see `sigpending(2)`)

Upon successful completion, each of the functions in the `exec` family marks for update the `st_atime` field of the file, unless the file is on a read-only file

system. Should the function succeed, the process image file is considered to have been opened by the `open(2)` system called. The corresponding `close()` is considered to occur at a time after this open, but before process termination or successful completion of a subsequent call to one of the functions in the `exec` family.

RETURN VALUES

If a function in the `exec` family returns to the calling process, an error has occurred; the return value is `-1` and `errno` is set to indicate the error.

ERRORS

Each of the functions in the `exec` family will fail if:

- E2BIG** The number of bytes in the new process's argument list is greater than the system-imposed limit of `ARG_MAX` bytes. The argument list limit is sum of the size of the argument list plus the size of the environment's exported shell variables.
- EACCES** Search permission is denied for a directory listed in the new process file's path prefix; the new process file is not an ordinary file; or the new process file mode denies execute permission.
- EAGAIN** Total amount of system memory available when reading using raw I/O is temporarily insufficient.
- EFAULT** An argument points to an illegal address.
- EINTR** A signal was caught during the execution of one of the functions in the `exec` family.
- ELOOP** Too many symbolic links were encountered in translating *path* or *file*.
- ENAMETOOLONG** The length of the *file* or *path* argument exceeds `PATH_MAX`, or the length of a *file* or *path* component exceeds `{ NAME_MAX }` while `{ _POSIX_NO_TRUNC }` is in effect.
- ENOENT** One or more components of the new process path name of the file do not exist or is a null pathname.
- ENOEXEC** The function call is not an `execlp()` or `execvp()`, and the new process file has the appropriate access permission but an invalid magic number in its header.
- ENOLINK** The *path* argument points to a remote machine and the link to that machine is no longer active.
- ENOMEM** The new process requires more memory than is allowed by the limit imposed by `getrlimit()` (see `brk(2)`).

ENOTDIR A component of the new process path of the file prefix is not a directory.

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	execle() and execve() are Async-Signal-Safe

SEE ALSO

ksh(1), **ps(1)**, **sh(1)**, **alarm(2)**, **brk(2)**, **chmod(2)**, **exit(2)**, **fcntl(2)**, **fork(2)**, **getrlimit(2)**, **mmap(2)**, **nice(2)**, **priocntl(2)**, **profil(2)**, **ptrace(2)**, **semop(2)**, **shmop(2)**, **sigpending(2)**, **sigprocmask(2)**, **times(2)**, **umask(2)**, **lockf(3C)**, **signal(3C)**, **system(3S)**, **timer_create(3R)**, **a.out(4)**, **attributes(5)**, **environ(5)**, **standards(5)**

WARNINGS

If a program is **setuid** to a user ID other than the super-user, and the program 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, _exit – terminate process
SYNOPSIS	<pre>#include <stdlib.h> void exit(int <i>status</i>); #include <unistd.h> void _exit(int <i>status</i>);</pre>
DESCRIPTION	<p>The exit() function first calls all functions registered by atexit(3C), in the reverse order of their registration. Each function is called as many times as it was registered.</p> <p>If a function registered by a call to atexit(3C) fails to return, the remaining registered functions are not called and the rest of the exit() processing is not completed. If exit() is called more than once, the effects are undefined.</p> <p>The exit() function then flushes all output streams, closes all open streams, and removes all files created by tmpfile(3S).</p> <p>The _exit() and exit() functions terminate the calling process with the following consequences:</p> <ul style="list-style-type: none"> ■ All of the file descriptors, directory streams, conversion descriptors and message catalogue descriptors open in the calling process are closed. ■ If the parent process of the calling process is executing a wait(2), wait3(3C), waitid(2) or waitpid(2), and has neither set its SA_NOCLDWAIT flag nor set SIGCHLD to SIG_IGN, it is notified of the calling process's termination and the low-order eight bits (that is, bits 0377) of <i>status</i> are made available to it. If the parent is not waiting, the child's status will be made available to it when the parent subsequently executes wait(2), wait3(3C), waitid(2) or waitpid(2). ■ If the parent process of the calling process is not executing a wait(2), wait3(3C), waitid(2) or waitpid(2), and has not set its SA_NOCLDWAIT flag, or set SIGCHLD to SIG_IGN, the calling process is transformed into a <i>zombie process</i>. A <i>zombie process</i> is an inactive process and it will be deleted at some later time when its parent process executes wait(2), wait3(3C), waitid(2) or waitpid(2). A zombie process only occupies a slot in the process table; it has no other space allocated either in user or kernel space. The process table slot that it occupies is partially overlaid with time accounting information (see <code><sys/proc.h></code>) to be used by the times(2) function. ■ Termination of a process does not directly terminate its children. The sending of a SIGHUP signal as described below indirectly terminates children in some circumstances.

- A SIGCHLD will be sent to the parent process.
- The parent process ID of all of the calling process's existing child processes and zombie processes is set to 1. That is, these processes are inherited by the initialization process (see `intro(2)`).
- Each mapped memory object is unmapped.
- Each attached shared-memory segment is detached and the value of `shm_nattch` (see `shmget(2)`) in the data structure associated with its shared memory ID is decremented by 1.
- For each semaphore for which the calling process has set a `semadj` value (see `semop(2)`), that value is added to the `semval` of the specified semaphore.
- If the process is a controlling process, the SIGHUP signal will be sent to each process in the foreground process group of the controlling terminal belonging to the calling process.
- If the process is a controlling process, the controlling terminal associated with the session is disassociated from the session, allowing it to be acquired by a new controlling process.
- If the exit of the process causes a process group to become orphaned, and if any member of the newly-orphaned process group is stopped, then a SIGHUP signal followed by a SIGCONT signal will be sent to each process in the newly-orphaned process group.
- If the parent process has set its `SA_NOCLDWAIT` flag, or set SIGCHLD to `SIG_IGN`, the status will be discarded, and the lifetime of the calling process will end immediately.
- If the process has process, text or data locks, an UNLOCK is performed (see `plock(3C)` and `memcntl(2)`).
- All open named semaphores in the process are closed as if by appropriate calls to `sem_close(3R)`. All open message queues in the process are closed as if by appropriate calls to `mq_close(3R)`. Any outstanding asynchronous I/O operations may be cancelled.
- An accounting record is written on the accounting file if the system's accounting routine is enabled (see `acct(2)`).

RETURN VALUES

These functions do not return.

ERRORS

No errors are defined.

USAGE

Normally applications should use `exit()` rather than `_exit()`.

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	<code>_exit()</code> is Async-Signal Safe

SEE ALSO

intro(2), **acct(2)**, **close(2)**, **memcntl(2)**, **semop(2)**, **shmget(2)**, **sigaction(2)**, **times(2)**, **wait(2)**, **waitid(2)**, **waitpid(2)**, **atexit(3C)**, **fclose(3S)**, **mq_close(3R)**, **plock(3C)**, **tmpfile(3S)**, **wait3(3C)**, **attributes(5)**, **signal(5)**

NAME	fcntl – file control										
SYNOPSIS	<pre>#include <sys/types.h> #include <unistd.h> #include <fcntl.h> int fcntl(int <i>fildev</i>, int <i>cmd</i>, /* <i>arg</i> */ ...);</pre>										
DESCRIPTION	<p>The fcntl() function provides for control over open files. The <i>fildev</i> argument is an open file descriptor.</p> <p>The fcntl() function may take a third argument, <i>arg</i>, whose data type, value and use depend upon the value of <i>cmd</i>. The <i>cmd</i> argument specifies the operation to be performed by fcntl().</p> <p>The available values for <i>cmd</i> are defined in the header <code><fcntl.h></code>, which include:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 10px;">F_DUPFD</td> <td>Return a new file descriptor which is the lowest numbered available (that is, not already open) file descriptor greater than or equal to the third argument, <i>arg</i>, taken as an integer of type <code>int</code>. The new file descriptor refers to the same open file description as the original file descriptor, and shares any locks. The <code>FD_CLOEXEC</code> flag associated with the new file descriptor is cleared to keep the file open across calls to one of the exec(2) functions.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 10px;">F_DUP2FD</td> <td>Similar to <code>F_DUPFD</code>, but always returns <i>arg</i>. <code>F_DUP2FD</code> closes <i>arg</i> if it is open and not equal to <i>fildev</i>. <code>F_DUP2FD</code> is equivalent to <code>dup2(fildev, arg)</code>.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 10px;">F_GETFD</td> <td>Get the file descriptor flags defined in <code><fcntl.h></code> that are associated with the file descriptor <i>fildev</i>. File descriptor flags are associated with a single file descriptor and do not affect other file descriptors that refer to the same file.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 10px;">F_SETFD</td> <td>Set the file descriptor flags defined in <code><fcntl.h></code>, that are associated with <i>fildev</i>, to the third argument, <i>arg</i>, taken as type <code>int</code>. If the <code>FD_CLOEXEC</code> flag in the third argument is 0, the file will remain open across the exec() functions; otherwise the file will be closed upon successful execution of one of the exec() functions.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 10px;">F_GETFL</td> <td>Get the file status flags and file access modes, defined in <code><fcntl.h></code>, for the file description associated with <i>fildev</i>. The file access modes can be extracted from the return value using the mask <code>O_ACCMODE</code>, which is defined in <code><fcntl.h></code>. File status flags and file access modes are</td> </tr> </table>	F_DUPFD	Return a new file descriptor which is the lowest numbered available (that is, not already open) file descriptor greater than or equal to the third argument, <i>arg</i> , taken as an integer of type <code>int</code> . The new file descriptor refers to the same open file description as the original file descriptor, and shares any locks. The <code>FD_CLOEXEC</code> flag associated with the new file descriptor is cleared to keep the file open across calls to one of the exec(2) functions.	F_DUP2FD	Similar to <code>F_DUPFD</code> , but always returns <i>arg</i> . <code>F_DUP2FD</code> closes <i>arg</i> if it is open and not equal to <i>fildev</i> . <code>F_DUP2FD</code> is equivalent to <code>dup2(fildev, arg)</code> .	F_GETFD	Get the file descriptor flags defined in <code><fcntl.h></code> that are associated with the file descriptor <i>fildev</i> . File descriptor flags are associated with a single file descriptor and do not affect other file descriptors that refer to the same file.	F_SETFD	Set the file descriptor flags defined in <code><fcntl.h></code> , that are associated with <i>fildev</i> , to the third argument, <i>arg</i> , taken as type <code>int</code> . If the <code>FD_CLOEXEC</code> flag in the third argument is 0, the file will remain open across the exec() functions; otherwise the file will be closed upon successful execution of one of the exec() functions.	F_GETFL	Get the file status flags and file access modes, defined in <code><fcntl.h></code> , for the file description associated with <i>fildev</i> . The file access modes can be extracted from the return value using the mask <code>O_ACCMODE</code> , which is defined in <code><fcntl.h></code> . File status flags and file access modes are
F_DUPFD	Return a new file descriptor which is the lowest numbered available (that is, not already open) file descriptor greater than or equal to the third argument, <i>arg</i> , taken as an integer of type <code>int</code> . The new file descriptor refers to the same open file description as the original file descriptor, and shares any locks. The <code>FD_CLOEXEC</code> flag associated with the new file descriptor is cleared to keep the file open across calls to one of the exec(2) functions.										
F_DUP2FD	Similar to <code>F_DUPFD</code> , but always returns <i>arg</i> . <code>F_DUP2FD</code> closes <i>arg</i> if it is open and not equal to <i>fildev</i> . <code>F_DUP2FD</code> is equivalent to <code>dup2(fildev, arg)</code> .										
F_GETFD	Get the file descriptor flags defined in <code><fcntl.h></code> that are associated with the file descriptor <i>fildev</i> . File descriptor flags are associated with a single file descriptor and do not affect other file descriptors that refer to the same file.										
F_SETFD	Set the file descriptor flags defined in <code><fcntl.h></code> , that are associated with <i>fildev</i> , to the third argument, <i>arg</i> , taken as type <code>int</code> . If the <code>FD_CLOEXEC</code> flag in the third argument is 0, the file will remain open across the exec() functions; otherwise the file will be closed upon successful execution of one of the exec() functions.										
F_GETFL	Get the file status flags and file access modes, defined in <code><fcntl.h></code> , for the file description associated with <i>fildev</i> . The file access modes can be extracted from the return value using the mask <code>O_ACCMODE</code> , which is defined in <code><fcntl.h></code> . File status flags and file access modes are										

associated with the file description and do not affect other file descriptors that refer to the same file with different open file descriptions.

F_SETFL Set the file status flags, defined in `<fcntl.h>`, for the file description associated with *fil* from the corresponding bits in the third argument, *arg*, taken as type `int`. Bits corresponding to the file access mode and the *oflag* values that are set in *arg* are ignored. If any bits in *arg* other than those mentioned here are changed by the application, the result is unspecified.

F_GETOWN If *fil* refers to a socket, get the process or process group ID specified to receive `SIGURG` signals when out-of-band data is available. Positive values indicate a process ID; negative values, other than `-1`, indicate a process group ID. If *fil* does not refer to a socket, the results are unspecified.

F_SETOWN If *fil* refers to a socket, set the process or process group ID specified to receive `SIGURG` signals when out-of-band data is available, using the value of the third argument, *arg*, taken as type `int`. Positive values indicate a process ID; negative values, other than `-1`, indicate a process group ID. If *fil* does not refer to a socket, the results are unspecified.

F_FREESP Free storage space associated with a section of the ordinary file *fil*. The section is specified by a variable of data type `struct flock` pointed to by *arg*. The data type `struct flock` is defined in the `<fcntl.h>` header (see `fcntl(5)`) and is described below. Note that all file systems might not support all possible variations of `F_FREESP` arguments. In particular, many file systems allow space to be freed only at the end of a file.

The following commands are available for advisory record locking. Record locking is supported for regular files, and may be supported for other files.

F_GETLK Get the first lock which blocks the lock description pointed to by the third argument, *arg*, taken as a pointer to type `struct flock`, defined in `<fcntl.h>`. The information retrieved overwrites the information passed to `fcntl()` in the structure `flock`. If no lock is found that would prevent this lock from being created, then the structure will be left unchanged except for the lock type which will be set to `F_UNLCK`.

F_GETLK64	Equivalent to F_GETLK, but takes a struct flock64 argument rather than a struct flock argument.
F_SETLK	Set or clear a file segment lock according to the lock description pointed to by the third argument, <i>arg</i> , taken as a pointer to type struct flock, defined in <fcntl.h>. F_SETLK is used to establish shared (or read) locks (F_RDLCK) or exclusive (or write) locks (F_WRLCK), as well as to remove either type of lock (F_UNLCK). F_RDLCK, F_WRLCK and F_UNLCK are defined in <fcntl.h>. If a shared or exclusive lock cannot be set, fcntl() will return immediately with a return value of -1.
F_SETLK64	Equivalent to F_SETLK, but takes a struct flock64 argument rather than a struct flock argument.
F_SETLKW	This command is the same as F_SETLK except that if a shared or exclusive lock is blocked by other locks, the process will wait until the request can be satisfied. If a signal that is to be caught is received while fcntl() is waiting for a region, fcntl() will be interrupted. Upon return from the process' signal handler, fcntl() will return -1 with errno set to EINTR, and the lock operation will not be done.
F_SETLKW64	Equivalent to F_SETLKW, but takes a struct flock64 argument rather than a struct flock argument.

When a shared lock is set on a segment of a file, other processes will be able to set shared locks on that segment or a portion of it. A shared lock prevents any other process from setting an exclusive lock on any portion of the protected area. A request for a shared lock will fail if the file descriptor was not opened with read access.

An exclusive lock will prevent any other process from setting a shared lock or an exclusive lock on any portion of the protected area. A request for an exclusive lock will fail if the file descriptor was not opened with write access.

The flock structure contains at least the following elements:

```

short  l_type;      /* lock operation type */
short  l_whence;    /* lock base indicator */
off_t  l_start;     /* starting offset from base */
off_t  l_len;       /* lock length; l_len == 0 means
                    until end of file */
long   l_sysid;     /* system ID running process holding lock */
pid_t  l_pid;       /* process ID of process holding lock */

```

The value of `l_whence` is `SEEK_SET`, `SEEK_CUR`, or `SEEK_END`, to indicate that the relative offset `l_start` bytes will be measured from the start of the file, current position or end of the file, respectively. The value of `l_len` is the number of consecutive bytes to be locked. The value of `l_len` may be negative (where the definition of `off_t` permits negative values of `l_len`). After a successful `F_GETLK` or `F_GETLK64` request, that is, one in which a lock was found, the value of `l_whence` will be `SEEK_SET`.

The `l_pid` and `l_sysid` fields are used only with `F_GETLK` or `F_GETLK64` to return the process ID of the process holding a blocking lock and to indicate which system is running that process.

If `l_len` is positive, the area affected starts at `l_start` and ends at `l_start + l_len - 1`. If `l_len` is negative, the area affected starts at `l_start + l_len` and ends at `l_start - 1`. Locks may start and extend beyond the current end of a file, but must not be negative relative to the beginning of the file. A lock will be set to extend to the largest possible value of the file offset for that file by setting `l_len` to 0. If such a lock also has `l_start` set to 0 and `l_whence` is set to `SEEK_SET`, the whole file will be locked.

If a process has an existing lock in which `l_len` is 0 and which includes the last byte of the requested segment, and an unlock (`F_UNLCK`) request is made in which `l_len` is non-zero and the offset of the last byte of the requested segment is the maximum value for an object of type `off_t`, then the `F_UNLCK` request will be treated as a request to unlock from the start of the requested segment with an `l_len` equal to 0. Otherwise, the request will attempt to unlock only the requested segment.

There will be at most one type of lock set for each byte in the file. Before a successful return from an `F_SETLK`, `F_SETLK64`, `F_SETLKW`, or `F_SETLKW64` request when the calling process has previously existing locks on bytes in the region specified by the request, the previous lock type for each byte in the specified region will be replaced by the new lock type. As specified above under the descriptions of shared locks and exclusive locks, an `F_SETLK`, `F_SETLK64`, `F_SETLKW`, or `F_SETLKW64` request will (respectively) fail or block when another process has existing locks on bytes in the specified region and the type of any of those locks conflicts with the type specified in the request.

All locks associated with a file for a given process are removed when a file descriptor for that file is closed by that process or the process holding that file descriptor terminates. Locks are not inherited by a child process created using `fork(2)`.

A potential for deadlock occurs if a process controlling a locked region is put to sleep by attempting to lock another process' locked region. If the system

detects that sleeping until a locked region is unlocked would cause a deadlock, **fcntl()** will fail with an EDEADLK error.

The following values for *cmd* are used for file share reservations. A share reservation is placed on an entire file to allow cooperating processes to control access to the file.

F_SHARE Sets a share reservation on a file with the specified access mode and designates which types of access to deny.

F_UNSHARE Remove an existing share reservation.

File share reservations are an advisory form of access control among cooperating processes, on both local and remote machines. They are most often used by DOS or Windows emulators and DOS based NFS clients. However, native UNIX versions of DOS or Windows applications may also choose to use this form of access control.

A share reservation is described by an *fshare* structure defined in `<sys/fcntl.h>`, which is included in `<fcntl.h>` as follows:

```
typedef struct fshare {
    short    f_access;
    short    f_deny;
    long     f_id;
} fshare_t;
```

A share reservation specifies the type of access, *f_access*, to be requested on the open file descriptor. If access is granted, it further specifies what type of access to deny other processes, *f_deny*. A single process on the same file may hold multiple non-conflicting reservations by specifying an identifier, *f_id*, unique to the process, with each request.

An **F_UNSHARE** request releases the reservation with the specified *f_id*. The *f_access* and *f_deny* fields are ignored.

Valid *f_access* values are:

F_RDACC Set a file share reservation for read-only access.

F_WRACC Set a file share reservation for write-only access.

F_RWACC Set a file share reservation for read and write access.

Valid *f_deny* values are:

F_COMPAT Set a file share reservation to compatibility mode.

F_RDDNY Set a file share reservation to deny read access to other processes.

RETURN VALUES

F_WRDNY	Set a file share reservation to deny write access to other processes.
F_RWDNY	Set a file share reservation to deny read and write access to other processes.
F_NODNY	Do not deny read or write access to any other process.
Upon successful completion, the value returned depends on <i>cmd</i> as follows:	
F_DUPFD	A new file descriptor.
F_GETFD	Value of flags defined in <code><fcntl.h></code> . The return value will not be negative.
F_SETFD	Value other than <code>-1</code> .
F_GETFL	Value of file status flags and access modes. The return value will not be negative.
F_SETFL	Value other than <code>-1</code> .
F_GETOWN	Value of the socket owner process or process group; this will not be <code>-1</code> .
F_SETOWN	Value other than <code>-1</code> .
F_FREESP	Value of <code>0</code> .
F_GETLK	Value other than <code>-1</code> .
F_GETLK64	Value other than <code>-1</code> .
F_SETLK	Value other than <code>-1</code> .
F_SETLK64	Value other than <code>-1</code> .
F_SETLKW	Value other than <code>-1</code> .
F_SETLKW64	Value other than <code>-1</code> .
F_SHARE	Value other than <code>-1</code> .
F_UNSHARE	Value other than <code>-1</code> .
Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.	

ERRORS

The `fcntl()` function will fail if:

EAGAIN

The *cmd* argument is `F_SETLK` or `F_SETLK64`, the type of lock (*l_type*) is a shared (`F_RDLCK`) or exclusive (`F_WRLCK`) lock, and the segment of a file to be locked is already exclusive-locked by another process; or the type is an exclusive lock and some portion of the segment of a file to be locked is already shared-locked or exclusive-locked by another process.

The *cmd* argument is `F_FREESP`, the file exists, mandatory file/record locking is set, and there are outstanding record locks on the file; or the *cmd* argument is `F_SETLK`, `F_SETLK64`, `F_SETLKW`, or `F_SETLKW64`, mandatory file/record locking is set, and the file is currently being mapped to virtual memory using `mmap(2)`.

The *cmd* argument is `F_SHARE` and *f_access* conflicts with an existing *f_deny* share reservation.

EBADF

The *fildev* argument is not a valid open file descriptor; or the *cmd* argument is `F_SETLK`, `F_SETLK64`, `F_SETLKW`, or `F_SETLKW64`, the type of lock, *l_type*, is a shared lock (`F_RDLCK`), and *fildev* is not a valid file descriptor open for reading; or the type of lock *l_type* is an exclusive lock (`F_WRLCK`) and *fildev* is not a valid file descriptor open for writing.

The *cmd* argument is `F_FREESP` and *fildev* is not a valid file descriptor open for writing.

The *cmd* argument is `F_DUP2FD`, and *arg* is negative or is not less than the current resource limit for `RLIMIT_NOFILE`.

The *cmd* argument is `F_SHARE`, the *f_access* share reservation is for write access, and *fildev* is not a valid file descriptor open for writing.

The *cmd* argument is `F_SHARE`, the *f_access* share reservation is for read access, and *fildev* is not a valid file descriptor open for reading.

EFAULT

The *cmd* argument is `F_GETLK`, `F_GETLK64`, `F_SETLK`, `F_SETLK64`, `F_SETLKW`, `F_SETLKW64`, or `F_FREESP` and the *arg* argument points to an illegal address.

The *cmd* argument is `F_SHARE` or `F_UNSHARE` and *arg* points to an illegal address.

EINTR	The <i>cmd</i> argument is <code>F_SETLKW</code> or <code>F_SETLKW64</code> and the function was interrupted by a signal.
EINVAL	The <i>cmd</i> argument is invalid; or the <i>cmd</i> argument is <code>F_DUPFD</code> and <i>arg</i> is negative or greater than or equal to <code>OPEN_MAX</code> ; or the <i>cmd</i> argument is <code>F_GETLK</code> , <code>F_GETLK64</code> , <code>F_SETLK</code> , <code>F_SETLK64</code> , <code>F_SETLKW</code> , or <code>F_SETLKW64</code> and the data pointed to by <i>arg</i> is not valid; or <i>fildev</i> refers to a file that does not support locking. The <i>cmd</i> argument is <code>F_UNSHARE</code> and a reservation with this <i>f_id</i> for this process does not exist.
EIO	An I/O error occurred while reading from or writing to the file system.
EMFILE	The <i>cmd</i> argument is <code>F_DUPFD</code> and either <code>OPEN_MAX</code> file descriptors are currently open in the calling process, or no file descriptors greater than or equal to <i>arg</i> are available.
ENOLCK	The <i>cmd</i> argument is <code>F_SETLK</code> , <code>F_SETLK64</code> , <code>F_SETLKW</code> , or <code>F_SETLKW64</code> and satisfying the lock or unlock request would result in the number of locked regions in the system exceeding a system-imposed limit.
ENOLINK	Either the <i>fildev</i> argument is on a remote machine and the link to that machine is no longer active; or the <i>cmd</i> argument is <code>F_FREESP</code> , the file is on a remote machine, and the link to that machine is no longer active.
EOVERFLOW	One of the values to be returned cannot be represented correctly. The <i>cmd</i> argument is <code>F_GETLK</code> , <code>F_SETLK</code> , or <code>F_SETLKW</code> and the smallest or, if <i>l_len</i> is non-zero, the largest, offset of any byte in the requested segment cannot be represented correctly in an object of type <code>off_t</code> . The <i>cmd</i> argument is <code>F_GETLK64</code> , <code>F_SETLK64</code> , or <code>F_SETLKW64</code> and the smallest or, if <i>l_len</i> is non-zero, the largest, offset of any byte in the requested segment cannot be represented correctly in an object of type <code>off64_t</code> .
The <code>fcntl()</code> function may fail if:	
EAGAIN	The <i>cmd</i> argument is <code>F_SETLK</code> , <code>F_SETLK64</code> , <code>F_SETLKW</code> , or <code>F_SETLKW64</code> , and the file is currently being mapped to virtual memory using <code>mmap(2)</code> .

EDEADLK The *cmd* argument is `F_SETLKW` or `F_SETLKW64`, the lock is blocked by some lock from another process and putting the calling process to sleep, waiting for that lock to become free would cause a deadlock.

The *cmd* argument is `F_FREESP`, mandatory record locking is enabled, `O_NDELAY` and `O_NONBLOCK` are clear and a deadlock condition was detected.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal Safe

SEE ALSO

`lockd(1M)`, `chmod(2)`, `close(2)`, `creat(2)`, `dup(2)`, `exec(2)`, `fork(2)`, `mmap(2)`, `open(2)`, `pipe(2)`, `read(2)`, `sigaction(2)`, `write(2)`, `dup2(3C)`, `attributes(5)`, `fcntl(5)`

System Interface Guide

NOTES

In the past, the variable `errno` was set to `EACCES` rather than `EAGAIN` when a section of a file is already locked by another process. Therefore, portable application programs should expect and test for either value.

Advisory locks allow cooperating processes to perform consistent operations on files, but do not guarantee exclusive access. Files can be accessed without advisory locks, but inconsistencies may result. The network share locking protocol does not support the `f_deny` value of `F_COMPAT`. For network file systems, if `f_access` is `F_RDACC`, `f_deny` is mapped to `F_RDDNY`. Otherwise, it is mapped to `F_RWDNY`.

If the file server crashes and has to be rebooted, the lock manager (see `lockd(1M)`) 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.

NAME	fork, fork1 – create a new process
SYNOPSIS	<pre>#include <sys/types.h> #include <unistd.h> pid_t fork(void); pid_t fork1(void);</pre>
DESCRIPTION	<p>The fork() and fork1() functions create a new process. The new process (child process) is an exact copy of the calling process (parent process). The child process inherits the following attributes from the parent process:</p> <ul style="list-style-type: none">■ real user ID, real group ID, effective user ID, effective group ID■ environment■ open file descriptors■ close-on-exec flags (see exec(2))■ signal handling settings (that is, SIG_DFL , SIG_IGN , SIG_HOLD , function address)■ supplementary group IDs■ set-user-ID mode bit■ set-group-ID mode bit■ profiling on/off status■ nice value (see nice(2))■ scheduler class (see pricon(1))■ all attached shared memory segments (see shmop(2))■ process group ID – memory mappings (see mmap(2))■ session ID (see exit(2))■ current working directory■ root directory■ file mode creation mask (see umask(2))■ resource limits (see getrlimit(2))■ controlling terminal■ saved user ID and group ID

Scheduling priority and any per-process scheduling parameters that are specific to a given scheduling class may or may not be inherited according to the policy of that particular class (see `pricntl(2)`). The child process differs from the parent process in the following ways:

- The child process has a unique process ID which does not match any active process group ID.
- The child process has a different parent process ID (that is, the process ID of the parent process).
- The child process has its own copy of the parent's file descriptors and directory streams. Each of the child's file descriptors shares a common file pointer with the corresponding file descriptor of the parent.
- Each shared memory segment remains attached and the value of `shm_nattach` is incremented by 1.
- All `semadj` values are cleared (see `semop(2)`).
- Process locks, text locks, data locks, and other memory locks are not inherited by the child (see `plock(3C)` and `memcntl(2)`).
- The child process's `tms` structure is cleared: `tms_utime`, `stime`, `cutime`, and `cstime` are set to 0 (see `times(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 set of signals pending for the child process is initialized to the empty set.
- Timers created by `timer_create(3R)` are not inherited by the child process.
- No asynchronous input or asynchronous output operations are inherited by the child.

Record locks set by the parent process are not inherited by the child process (see `fcntl(2)`).

MT fork() Solaris Threads

The following are the `fork()` semantics in programs that use the Solaris threads API rather than the POSIX threads (see `standards(5)`) API (programs linked with `-lthread` but not `-lpthread`):

The `fork()` function duplicates all the threads (see `thr_create(3T)`) and LWPs in the parent process in the child process. The `fork1()` function duplicates only the calling thread (LWP) in the child process.

POSIX Threads

The following are the **fork()** semantics in programs that use the POSIX threads API rather than the Solaris threads API (programs linked with `-lpthread`, whether or not linked with `-lthread`):

The call to **fork()** is like a call to **fork1()**, which replicates only the calling thread. There is no call that forks a child with all threads and LWPs duplicated in the child.

Note that if a program is linked with both libraries (`-lthread` and `-lpthread`), the POSIX semantic of **fork()** prevails.

Fork-safety

If **fork1()** is called in a Solaris thread program or **fork()** is called in a POSIX thread program, and the child does more than just call **exec()**, there is a possibility of deadlocking in the child. To ensure that the application is safe with respect to this deadlock, it should use **pthread_atfork(3T)**. Should there be any outstanding mutexes throughout the process, the application should call **pthread_atfork(3T)**, to wait for and acquire those mutexes, prior to calling **fork()**. (See **attributes(5)** "MT-Level of Libraries")

RETURN VALUES

Upon successful completion, **fork()** and **fork1()** return 0 to the child process and return the process ID of the child process to the parent process. Otherwise, `(pid_t)-1` is returned to the parent process, no child process is created, and `errno` is set to indicate the error.

ERRORS

The **fork()** function will fail if:

EAGAIN The system-imposed limit on the total number of processes under execution by a single user has been exceeded; or the total amount of system memory available is temporarily insufficient to duplicate this process.

ENOMEM There is not enough swap space.

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	<code>fork()</code> is Async-Signal-Safe

SEE ALSO

alarm(2), **exec(2)**, **exit(2)**, **fcntl(2)**, **getitimer(2)**, **getrlimit(2)**, **memcntl(2)**, **mmap(2)**, **nice(2)**, **prctl(2)**, **ptrace(2)**, **semop(2)**, **shmop(2)**, **times(2)**, **umask(2)**, **wait(2)**, **exit(3C)**, **plock(3C)**, **pthread_atfork(3T)**, **signal(3C)**, **system(3S)**, **thr_create(3T)**, **timer_create(3R)**, **attributes(5)**, **standards(5)**

NOTES

Be careful to call `_exit()` rather than `exit(3C)` if you cannot `execve()`, since `exit(3C)` will flush and close standard I/O channels, and thereby corrupt the parent processes standard I/O data structures. Using `exit(3C)` will flush buffered data twice. See `exit(2)`.

When calling `fork1()` the thread (or LWP) in the child must not depend on any resources that are held by threads (or LWPs) that no longer exist in the child. In particular, locks held by these threads (or LWPs) will not be released.

In a multithreaded process, `fork()` or `fork1()` can cause blocking system calls to be interrupted and return with an error of `EINTR`.

The `fork()` and `fork1()` functions suspend all threads in the process before proceeding. Threads which are executing in the kernel and are in an uninterruptible wait cannot be suspended immediately; and therefore, cause a delay before `fork()` and `fork1()` can complete. During this delay, all other threads will have already been suspended, and so the process will appear "hung."

NAME | fpathconf, pathconf – get configurable pathname variables

SYNOPSIS | #include <unistd.h>

long int **fpathconf**(int *filde*, int *name*);

long int **pathconf**(const char * *path*, int *name*);

DESCRIPTION

The **fpathconf()** and **pathconf()** functions provide a method for the application to determine the current value of a configurable limit or option I (variable)that is associated with a file or directory.

For **pathconf()** , the *path* argument points to the pathname of a file or directory.

For **fpathconf()** , the *filde* argument is an open file descriptor.

The *name* argument represents the variable to be queried relative to that file or directory. The variables in the following table come from <limits.h> or <unistd.h> and the symbolic constants, defined in <unistd.h> , are the corresponding values used for *name*:

Variable	Value of <i>name</i>	Notes
FILESIZEBITS	_PC_FILESIZEBITS	3,4
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
_POSIX_NO_TRUNC	_PC_NO_TRUNC	3,4
_POSIX_VDISABLE	_PC_VDISABLE	2
_POSIX_ASYNC_IO	_PC_ASYNC_IO	8
_POSIX_PRIO_IO	_PC_PRIO_IO	8
_POSIX_SYNC_IO	_PC_SYNC_IO	8

Notes:

1. If *path* or *filde* refers to a directory, the value returned applies to the directory itself.

2. If *path* or *fildev* does not refer to a terminal file, it is unspecified whether an implementation supports an association of the variable name with the specified file.
3. If *path* or *fildev* refers to a directory, the value returned applies to filenames within the directory.
4. If *path* or *fildev* does not refer to a directory, it is unspecified whether an implementation supports an association of the variable name with the specified file.
5. If *path* or *fildev* 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 *fildev* refers to a pipe or FIFO, the value returned applies to the referenced object. If *path* or *fildev* refers to a directory, the value returned applies to any FIFO that exists or can be created within the directory. If *path* or *fildev* refers to any other type of file, it is unspecified whether an implementation supports an association of the variable name with the specified file.
7. If *path* or *fildev* refers to a directory, the value returned applies to any files, other than directories, that exist or can be created within the directory.
8. If *path* or *fildev* refers to a directory, it is unspecified whether an implementation supports an association of the variable name with the specified file.

RETURN VALUES

If *name* is an invalid value, both **pathconf()** and **fpathconf()** return `-1` and `errno` is set to indicate the error.

If the variable corresponding to *name* has no limit for the *path* or file descriptor, both **pathconf()** and **fpathconf()** return `-1` without changing `errno`. If the implementation needs to use *path* to determine the value of *name* and the implementation does not support the association of *name* with the file specified by *path*, or if the process did not have appropriate privileges to query the *appropriate privileges* file specified by *path*, or *path* does not exist, **pathconf()** returns `-1` and `errno` is set to indicate the error.

If the implementation needs to use *fildev* to determine the value of *name* and the implementation does not support the association of *name* with the file specified by *fildev*, or if *fildev* is an invalid file descriptor, **fpathconf()** will return `-1` and `errno` is set to indicate the error.

Otherwise **pathconf()** or **fpathconf()** returns the current variable value for the file or directory without changing `errno`. The value returned will not be more restrictive than the corresponding value available to the application when it was compiled with the implementation's `<limits.h>` or `<unistd.h>`.

ERRORS

The **pathconf()** function will fail if:

EINVAL The value of *name* is not valid.

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

The **pathconf()** function may fail if:

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` or a pathname component is longer than `NAME_MAX` .

ENAMETOOLONG Pathname resolution of a symbolic link produced an intermediate result whose length exceeds `PATH_MAX` .

ENOENT A component of *path* does not name an existing file or *path* is an empty string.

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

The **fpathconf()** function will fail if:

EINVAL The value of *name* is not valid.

The **fpathconf()** function may fail if:

EBADF The *fildev* argument is not a valid file descriptor.

EINVAL The implementation does not support an association of the variable *name* with the specified file.

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	<code>pathconf()</code> is Async-Signal-Safe

SEE ALSO

sysconf(3C) , **limits(4)** , **attributes(5)** , **standards(5)**

NAME	getaudit, setaudit – get and set process audit information
SYNOPSIS	<pre>cc [<i>flag</i> ...] <i>file</i> ... -lbsm -lsocket -lnsl -lintl [<i>library</i> ...] #include <sys/param.h> #include <bsm/audit.h> int getaudit(struct auditinfo * <i>info</i>); int setaudit(struct auditinfo * <i>info</i>);</pre>
DESCRIPTION	<p>The getaudit() function gets the audit ID, the preselection mask, the terminal ID and the audit session ID for the current process.</p> <p>The setaudit() function sets the audit ID, the preselection mask, the terminal ID and the audit session ID for the current process.</p> <p>The <i>info</i> structure used to pass the process audit information contains the following members:</p> <pre>au_id_t ai_auid; /* audit user ID */ au_mask_t ai_mask; /* preselection mask */ au_tid_t ai_termid; /* terminal ID */ au_asid_t ai_asid; /* audit session ID */</pre>
RETURN VALUES	Upon successful completion, getaudit() and setaudit() return 0 . Otherwise, -1 is returned and <i>errno</i> is set to indicate the error.
ERRORS	<p>The getaudit() and setaudit() functions will fail if:</p> <p>EFAULT The <i>info</i> parameter points outside the process's allocated address space.</p>

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

USAGE Only processes with the effective user ID of the super-user may successfully execute these calls.

SEE ALSO **bsmconv(1M)** , **audit(2)**

NOTES The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See **bsmconv(1M)** for more information.

NAME	getaudit, setaudit – get and set user audit identity
SYNOPSIS	<pre>cc [<i>flag</i> ...] <i>file</i> ... -lbsm -lsocket -lnsl -lintl [<i>library</i> ...] #include <sys/param.h> #include <bsm/audit.h> int getaudit(au_id_t * <i>audit</i>); int setaudit(au_id_t * <i>audit</i>);</pre>
DESCRIPTION	<p>The getaudit() function 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.</p> <p>The setaudit() function sets the audit user ID for the current process.</p>
RETURN VALUES	<p>Upon successful completion, the getaudit() function returns the audit user ID of the current process on success. Otherwise, it returns <code>-1</code> and sets <code>errno</code> to indicate the error.</p> <p>Upon successful completion the setaudit() function returns <code>0</code> . Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.</p>
ERRORS	<p>The getaudit() and setaudit() functions will fail if:</p> <p>EFAULT The <i>audit</i> argument points to an invalid address.</p> <p>EPERM The process's effective user ID is not super-user.</p>
USAGE	Only the super-user may successfully execute these calls.

SEE ALSO | `bsmconv(1M)` , `audit(2)` , `getaudit(2)`

NOTES | The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See `bsmconv(1M)` for more information.

These system calls have been superseded by `getaudit()` and `setaudit()` .

NAME	getcontext, setcontext – get and set current user context
SYNOPSIS	<pre>#include <ucontext.h> int getcontext(ucontext_t * ucp); int setcontext(const ucontext_t * ucp);</pre>
DESCRIPTION	<p>The getcontext() function initializes the structure pointed to by <i>ucp</i> to the current user context of the calling process. The <code>ucontext_t</code> type that <i>ucp</i> points to defines the user context and includes the contents of the calling process' machine registers, the signal mask, and the current execution stack.</p> <p>The setcontext() function restores the user context pointed to by <i>ucp</i> . A successful call to setcontext() does not return; program execution resumes at the point specified by the <i>ucp</i> argument passed to setcontext() . The <i>ucp</i> argument should be created either by a prior call to getcontext() , or by being passed as an argument to a signal handler. If the <i>ucp</i> argument was created with getcontext() , program execution continues as if the corresponding call of getcontext() had just returned. If the <i>ucp</i> argument was created with makecontext(3C) , program execution continues with the function passed to makecontext(3C) . When that function returns, the process continues as if after a call to setcontext() with the <i>ucp</i> argument that was input to makecontext(3C) . If the <i>ucp</i> argument was passed to a signal handler, program execution continues with the program instruction following the instruction interrupted by the signal. If the <code>uc_link</code> member of the <code>ucontext_t</code> structure pointed to by the <i>ucp</i> argument is equal to 0, then this context is the main context, and the process will exit when this context returns. The effects of passing a <i>ucp</i> argument obtained from any other source are unspecified.</p>
RETURN VALUES	On successful completion, setcontext() does not return and getcontext() returns 0 . Otherwise, -1 is returned.
ERRORS	No errors are defined.
USAGE	<p>When a signal handler is executed, the current user context is saved and a new context is created. If the process leaves the signal handler via longjmp(3B) , then it is unspecified whether the context at the time of the corresponding setjmp(3B) call is restored and thus whether future calls to getcontext() will provide an accurate representation of the current context, since the context restored by longjmp(3B) may not contain all the information that setcontext() requires. Signal handlers should use siglongjmp(3C) or setcontext() instead.</p> <p>Portable applications should not modify or access the <code>uc_mcontext</code> member of <code>ucontext_t</code> . A portable application cannot assume that context includes</p>

any process-wide static data, possibly including `errno`. Users manipulating contexts should take care to handle these explicitly when required.

SEE ALSO

`sigaction(2)`, `sigaltstack(2)`, `sigprocmask(2)`, `bsd_signal(3C)`,
`makecontext(3C)`, `setjmp(3B)`, `sigsetjmp(3C)`, `ucontext(5)`

NAME	getdents – read directory entries and put in a file system independent format																
SYNOPSIS	<pre>#include <sys/dirent.h> int getdents(int <i>fildev</i>, struct dirent *<i>buf</i>, size_t <i>nbyte</i>);</pre>																
DESCRIPTION	<p>The getdents() function attempts to read <i>nbyte</i> bytes from the directory associated with the file descriptor <i>fildev</i> and to format them as file system independent directory entries in the buffer pointed to by <i>buf</i>. Since the file system independent directory entries are of variable lengths, in most cases the actual number of bytes returned will be less than <i>nbyte</i>. The file system independent directory entry is specified by the <code>dirent</code> structure. See dirent(4).</p> <p>On devices capable of seeking, getdents() starts at a position in the file given by the file pointer associated with <i>fildev</i>. Upon return from getdents(), the file pointer is incremented to point to the next directory entry.</p>																
RETURN VALUES	<p>Upon successful completion, a non-negative integer is returned indicating the number of bytes actually read. A return value of 0 indicates the end of the directory has been reached. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.</p>																
ERRORS	<p>The getdents() function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EBADF</td> <td>The <i>fildev</i> argument is not a valid file descriptor open for reading.</td> </tr> <tr> <td>EFAULT</td> <td>The <i>buf</i> argument points to an illegal address.</td> </tr> <tr> <td>EINVAL</td> <td>The <i>nbyte</i> argument is not large enough for one directory entry.</td> </tr> <tr> <td>EIO</td> <td>An I/O error occurred while accessing the file system.</td> </tr> <tr> <td>ENOENT</td> <td>The current file pointer for the directory is not located at a valid entry.</td> </tr> <tr> <td>ENOLINK</td> <td>The <i>fildev</i> argument points to a remote machine and the link to that machine is no longer active.</td> </tr> <tr> <td>ENOTDIR</td> <td>The <i>fildev</i> argument is not a directory.</td> </tr> <tr> <td>E_OVERFLOW</td> <td>The value of the <code>dirent</code> structure member <code>d_ino</code> or <code>d_off</code> cannot be represented in an <code>ino_t</code> or <code>off_t</code>.</td> </tr> </table>	EBADF	The <i>fildev</i> argument is not a valid file descriptor open for reading.	EFAULT	The <i>buf</i> argument points to an illegal address.	EINVAL	The <i>nbyte</i> argument is not large enough for one directory entry.	EIO	An I/O error occurred while accessing the file system.	ENOENT	The current file pointer for the directory is not located at a valid entry.	ENOLINK	The <i>fildev</i> argument points to a remote machine and the link to that machine is no longer active.	ENOTDIR	The <i>fildev</i> argument is not a directory.	E_OVERFLOW	The value of the <code>dirent</code> structure member <code>d_ino</code> or <code>d_off</code> cannot be represented in an <code>ino_t</code> or <code>off_t</code> .
EBADF	The <i>fildev</i> argument is not a valid file descriptor open for reading.																
EFAULT	The <i>buf</i> argument points to an illegal address.																
EINVAL	The <i>nbyte</i> argument is not large enough for one directory entry.																
EIO	An I/O error occurred while accessing the file system.																
ENOENT	The current file pointer for the directory is not located at a valid entry.																
ENOLINK	The <i>fildev</i> argument points to a remote machine and the link to that machine is no longer active.																
ENOTDIR	The <i>fildev</i> argument is not a directory.																
E_OVERFLOW	The value of the <code>dirent</code> structure member <code>d_ino</code> or <code>d_off</code> cannot be represented in an <code>ino_t</code> or <code>off_t</code> .																

USAGE The `getdents()` function was developed to implement the `readdir(3C)` function and should not be used for other purposes.

The `getdents()` function has a transitional interface for 64-bit file offsets. See `lf64(5)`.

SEE ALSO `readdir(3C)`, `dirent(4)`, `lf64(5)`

NAME	getgroups, setgroups – get or set supplementary group access list IDs				
SYNOPSIS	<pre>#include <unistd.h> int getgroups(int gidsetsize, gid_t * grouplist); int setgroups(int ngroups, const gid_t * grouplist);</pre>				
DESCRIPTION	<p>The getgroups() function gets the current supplemental group access list of the calling process and stores the result in the array of group IDs specified by <i>grouplist</i>. This array has <i>gidsetsize</i> entries and must be large enough to contain the entire list. This list cannot be larger than <code>NGROUPS_MAX</code>. If <i>gidsetsize</i> equals 0, getgroups() will return the number of groups to which the calling process belongs without modifying the array pointed to by <i>grouplist</i>.</p> <p>The setgroups() function sets the supplementary group access list of the calling process from the array of group IDs specified by <i>grouplist</i>. The number of entries is specified by <i>ngroups</i> and can not be greater than <code>NGROUPS_MAX</code>.</p>				
RETURN VALUES	Upon successful completion, getgroups() returns the number of supplementary group IDs set for the calling process and setgroups() returns 0. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.				
ERRORS	<p>The getgroups() and setgroups() functions will fail if:</p> <p>EFAULT A referenced part of the array pointed to by <i>grouplist</i> is an illegal address.</p> <p>The getgroups() function will fail if:</p> <p>EINVAL The value of <i>gidsetsize</i> is non-zero and less than the number of supplementary group IDs set for the calling process.</p> <p>The setgroups() function will fail if:</p> <p>EINVAL The value of <i>ngroups</i> is greater than <code>NGROUPS_MAX</code>.</p> <p>EPERM The effective user of the calling process is not super-user.</p>				
USAGE	This function may be invoked only by the super-user.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	groups(1) , chown(2) , getuid(2) , setuid(2) , getgrnam(3C) , initgroups(3C) , attributes(5)				

NAME	getitimer, setitimer – get or set value of interval timer
SYNOPSIS	<pre>#include <sys/time.h> int getitimer(int which, struct itimerval * value); int setitimer(int which, const struct itimerval * value, struct itimerval * ovalue);</pre>
DESCRIPTION	<p>The system provides each process with four interval timers, defined in <code>sys/time.h</code>. The getitimer() function stores the current value of the timer specified by <i>which</i> into the structure pointed to by <i>value</i>. The setitimer() function call sets the value of the timer specified by <i>which</i> to the value specified in the structure pointed to by <i>value</i>, and if <i>ovalue</i> is not <code>NULL</code>, stores the previous value of the timer in the structure pointed to by <i>ovalue</i>.</p> <p>A timer value is defined by the <code>itimerval</code> structure (see gettimeofday(3C) for the definition of <code>timeval</code>), which includes the following members:</p> <pre>struct timeval it_interval; /* timer interval */ struct timeval it_value; /* current value */</pre> <p>The <code>it_value</code> member indicates the time to the next timer expiration. The <code>it_interval</code> member specifies a value to be used in reloading <code>it_value</code> when the timer expires. Setting <code>it_value</code> to 0 disables a timer, regardless of the value of <code>it_interval</code>. Setting <code>it_interval</code> to 0 disables a timer after its next expiration (assuming <code>it_value</code> is non-zero).</p> <p>Time values smaller than the resolution of the system clock are rounded up to the resolution of the system clock, except for <code>ITIMER_REALPROF</code>, whose values are rounded up to the resolution of the profiling clock. The four timers are as follows:</p> <p><code>ITIMER_REAL</code></p> <p>Decrements in real time. A <code>SIGALRM</code> signal is delivered when this timer expires.</p> <p>In the current and previous releases, when <code>setitimer(ITIMER_REAL, ...)</code> is called in a multithreaded process linked with <code>-lthread</code> (Solaris threads) or <code>-lpthread</code> (POSIX threads; see standards(5)), the resulting <code>SIGALRM</code> is sent to the bound thread that called setitimer(); setitimer() has a per-thread semantic when called from a bound thread. This semantic will become obsolete in a future release. The semantic will move to a per-process semantic, with the resulting <code>SIGALRM</code> being sent to the process. The</p>

SIGALRM so generated is not maskable on this bound thread by any signal masking function, `pthread_sigmask(3T)`, `thr_sigsetmask(3T)`, or `sigprocmask(2)`. This is a bug that will not be fixed, since the per-thread semantic will be discontinued in the next release.

Also, calling this routine from an unbound thread is not guaranteed to work as in the case of bound threads. The resulting SIGALRM may be sent to some other thread (see `alarm(2)`). This is a bug and will not be fixed since the per-thread semantic is going to be discontinued.

Calling `setitimer(ITIMER_REAL, . . .)` from a process linked with `-lpthread` (POSIX threads) has the same behavior as Solaris threads described above, where a Solaris bound thread is the same as a POSIX thread in system scheduling scope and a Solaris unbound thread is the same as a POSIX thread in local scheduling scope.

Hence, for multithreaded (Solaris or POSIX) programs in the current and previous releases, the only reliable way to use the `ITIMER_REAL` flag is to call it from a bound thread which does not mask SIGALRM and to expect the SIGALRM to be delivered to this bound thread.

The current working of this flag is not being improved since some applications might depend on the current (slightly broken) semantic. When this semantic is discontinued in the future, it will be replaced with a per-process semantic, i.e. using this flag from any thread, bound or unbound, will result in the SIGALRM being sent to the process.

New MT applications should not use this flag, and should use `alarm(2)` instead.

ITIMER_VIRTUAL

Decrements in process virtual time. It runs only when the process is executing. A SIGVTALRM signal is delivered when it expires. (For multithreaded programs see the WARNINGS section below).

ITIMER_PROF

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 functions, programs using this timer must be prepared to restart interrupted functions. (For multithreaded programs see the WARNINGS section below).

ITIMER_REALPROF

Decrements in real time. It is designed to be used for real-time profiling of multithreaded programs. Each time the `ITIMER_REALPROF` timer expires, one counter in a set of counters maintained by the system for each lightweight process (lwp) is incremented. The counter corresponds to the state of the lwp at the time of the timer tick. All lwps executing in user mode when the timer expires are interrupted into system mode. When each lwp resumes execution in user mode, if any of the elements in its set of counters are non-zero, the `SIGPROF` signal is delivered to the lwp. The `SIGPROF` signal is delivered before any other signal except `SIGKILL`. This signal does not interrupt any in-progress function. A `siginfo` structure, defined in `<sys/siginfo.h>`, is associated with the delivery of the `SIGPROF` signal, and includes the following members:

```

si_tstamp;    /* high resolution timestamp */
si_syscall;   /* current syscall */
si_sysarg;    /* number of syscall arguments */
si_sysarg[];  /* actual syscall arguments */
si_fault;     /* last fault type */
si_faddr;     /* last fault address */
si_mstate[]; /* ticks in each microstate */

```

The enumeration of microstates (indices into `si_mstate`) is defined in `<sys/msacct.h>`. (For multithreaded programs see the `WARNINGS` section below).

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned and `errno` is set to indicate the error.

ERRORS

The `getitimer()` and `setitimer()` functions will fail if:

EINVAL The specified number of seconds is greater than 100,000,000, the number of microseconds is greater than or equal to 1,000,000, or the *which* argument is unrecognized.

The `setitimer()` function will fail if:

EACCES Either an unbound Solaris thread or a POSIX thread in local scheduling scope with a flag other than `ITIMER_REAL` called `setitimer()`.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `alarm(2)`, `sigprocmask(2)`, `gettimeofday(3C)`,
`pthread_attr_setscope(3T)`, `pthread_sigmask(3T)`, `sleep(3C)`,
`sysconf(3C)`, `attributes(5)`, `standards(5)`

WARNINGS All flags to `setitimer()` other than `ITIMER_REAL` behave as documented only with “bound” threads. Their ability to mask the signal works only with bound threads. If the call is made using one of these flags from an unbound thread, the system call returns `-1` and sets `errno` to `EACCES`.

These behaviors are the same for bound or unbound POSIX threads. A POSIX thread with system-wide scope, created by the call

```
pthread_attr_setscope(&attr, PTHREAD_SCOPE_SYSTEM);
```

is equivalent to a Solaris bound thread. A POSIX thread with local process scope, created by the call

```
pthread_attr_setscope(&attr, PTHREAD_SCOPE_PROCESS);
```

is equivalent to a Solaris unbound thread.

NOTES The microseconds field should not be equal to or greater than one second.

The `setitimer()` function is independent of the `alarm()` function.

Do not use `setitimer(ITIMER_REAL)` with the `sleep()` routine. A `sleep(3C)` call wipes out knowledge of the user signal handler for `SIGALRM`.

The `ITIMER_PROF` and `ITIMER_REALPROF` timers deliver the same signal and have different semantics. They cannot be used together.

The granularity of the resolution of alarm time is platform-dependent.

NAME	getmsg, getpmsg – get next message off a stream
SYNOPSIS	<pre>#include <stropts.h> int getmsg(int fildes, struct strbuf * ctlptr, struct strbuf * dataptr, int * flagsp); int getpmsg(int fildes, struct strbuf * ctlptr, struct strbuf * dataptr, int * bandp, int * flagsp);</pre>
DESCRIPTION	<p>The getmsg() function 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.</p> <p>The getpmsg() function behaved like getmsg(), but provides finer control over the priority of the messages received. Except where noted, all information pertaining to getmsg() also pertains to getpmsg().</p> <p>The <i>fildes</i> argument specifies a file descriptor referencing an open stream. The <i>ctlptr</i> and <i>dataptr</i> arguments each point to a <code>strbuf</code> structure, which contains the following members:</p> <pre>int maxlen; /* maximum buffer length */ int len; /* length of data */ char *buf; /* ptr to buffer */</pre> <p>The <i>buf</i> member points to a buffer into which the data or control information is to be placed, and the <i>maxlen</i> member indicates the maximum number of bytes this buffer can hold. On return, the <i>len</i> member contains the number of bytes of data or control information actually received; 0 if there is a zero-length control or data part; or -1 if no data or control information is present in the message. The <i>flagsp</i> argument should point to an integer that indicates the type of message the user is able to receive, as described below.</p> <p>The <i>ctlptr</i> argument holds the control part from the message and the <i>dataptr</i> argument holds the data part from the message. If <i>ctlptr</i> (or <i>dataptr</i>) is NULL or the <i>maxlen</i> member is -1, the control (or data) part of the message is not processed and is left on the stream head read queue. If <i>ctlptr</i> (or <i>dataptr</i>) is not NULL and there is no corresponding control (or data) part of the messages on the stream head read queue, <i>len</i> is set to -1. If the <i>maxlen</i> member 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 <i>len</i> is set to 0. If the <i>maxlen</i> member 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 <i>len</i> is set to 0. If the <i>maxlen</i></p>

member 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 .

By default, **getmsg()** processes the first available message on the stream head read queue. A user may, however, choose to retrieve only high priority messages by setting the integer pointed to by *flagsp* to *RS_HIPRI* . In this case, **getmsg()** processes the next message only if it is a high priority message.

If the integer pointed to by *flagsp* is 0, **getmsg()** retrieves any message available on the stream head read queue. In this case, on return, the integer pointed to by *flagsp* will be set to *RS_HIPRI* if a high priority message was retrieved, or to 0 otherwise.

For **getpmsg()** , the *flagsp* argument points to a bitmask with the following mutually-exclusive flags defined: *MSG_HIPRI* , *MSG_BAND* , and *MSG_ANY* . Like **getmsg()** , **getpmsg()** processes the first available message on the stream head read queue. A user may choose to retrieve only high-priority messages by setting the integer pointed to by *flagsp* to *MSG_HIPRI* and the integer pointed to by *bandp* to 0. In this case, **getpmsg()** will only process the next message if it is a high-priority message. In a similar manner, a user may choose to retrieve a message from a particular priority band by setting the integer pointed to by *flagsp* to *MSG_BAND* and the integer pointed to by *bandp* to the priority band of interest. In this case, **getpmsg()** will only process the next message if it is in a priority band equal to, or greater than, the integer pointed to by *bandp* , or if it is a high-priority message. If a user just wants to get the first message off the queue, the integer pointed to by *flagsp* should be set to *MSG_ANY* and the integer pointed to by *bandp* should be set to 0. On return, if the message retrieved was a high-priority message, the integer pointed to by *flagsp* will be set to *MSG_HIPRI* and the integer pointed to by *bandp* will be set to 0. Otherwise, the integer pointed to by *flagsp* will be set to *MSG_BAND* and the integer pointed to by *bandp* will be set to the priority band of the message.

If *O_NDELAY* and *O_NONBLOCK* are clear, **getmsg()** blocks until a message of the type specified by *flagsp* is available on the stream head read queue. If *O_NDELAY* or *O_NONBLOCK* has been set and a message of the specified type 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()** continues to operate normally, as described above, until the stream head read queue is empty. Thereafter, it returns 0 in the *len* member of *ctlptr* and *dataptr* .

RETURN VALUES

Upon successful completion, a non-negative value is returned. A return value of 0 indicates that a full message was read successfully. A return value of *MORECTL* indicates that more control information is waiting for retrieval. A

return value of MOREDATA indicates that more data are waiting for retrieval. A return value of MORECTL | MOREDATA indicates that both types of information remain. Subsequent **getmsg()** calls retrieve the remainder of the message. However, if a message of higher priority has been received by the stream head read queue, the next call to **getmsg()** will retrieve that higher priority message before retrieving the remainder of the previously received partial message.

ERRORS

The **getmsg()** and **getpmsg()** functions will fail if:

- EAGAIN** The `O_NDELAY` or `O_NONBLOCK` flag is set and no messages are available.
- EBADF** The *fildev* argument is not a valid file descriptor open for reading.
- EBADMSG** Queued message to be read is not valid for `getmsg` .
- EFAULT** The *ctlptr* , *dataptr* , *bandp* , or *flagsp* argument points to an illegal address.
- EINTR** A signal was caught during the execution of the `getmsg` function.
- EINVAL** An illegal value was specified in *flagsp* , or the stream referenced by *fildev* is linked under a multiplexor.
- ENOSTR** A stream is not associated with *fildev* .
The **getmsg()** function 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(2)` , `write(2)`
STREAMS Programming Guide

NAME	getpid, getpgrp, getppid, getpgid – get process, process group, and parent process IDs				
SYNOPSIS	<pre>#include <unistd.h> pid_t getpid(void); pid_t getpgrp(void); pid_t getppid(void); pid_t getpgid(pid_t pid);</pre>				
DESCRIPTION	<p>The getpid() function returns the process ID of the calling process.</p> <p>The getpgrp() function returns the process group ID of the calling process.</p> <p>The getppid() function returns the parent process ID of the calling process.</p> <p>The getpgid() function returns the process group ID of the process whose process ID is equal to <i>pid</i>, or the process group ID of the calling process, if <i>pid</i> is equal to 0.</p>				
RETURN VALUES	Upon successful completion, these functions return the process group ID. Otherwise, getpgid() returns <code>(pid_t)-1</code> and sets <code>errno</code> to indicate the error.				
ERRORS	<p>The getpgid() function will fail if:</p> <p>EPERM The process whose process ID is equal to <i>pid</i> is not in the same session as the calling process, and the implementation does not allow access to the process group ID of that process from the calling process.</p> <p>ESRCH There is no process with a process ID equal to <i>pid</i>.</p> <p>The getpgid() function may fail if:</p> <p>EINVAL The value of the <i>pid</i> argument is invalid.</p>				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	intro(2) , exec(2) , fork(2) , getsid(2) , setpgid(2) , setpgrp(2) , setsid(2) , signal(3C) , attributes(5)				

NAME	getrlimit, setrlimit – control maximum system resource consumption						
SYNOPSIS	<pre>#include <sys/resource.h> int getrlimit(int <i>resource</i>, struct rlimit * <i>rlp</i>); int setrlimit(int <i>resource</i>, const struct rlimit * <i>rlp</i>);</pre>						
DESCRIPTION	<p>Limits on the consumption of a variety of system resources by a process and each process it creates may be obtained with the getrlimit() and set with setrlimit() functions.</p> <p>Each call to either getrlimit() or setrlimit() identifies a specific resource to be operated upon as well as a resource limit. A resource limit is a pair of values: one specifying the current (soft) limit, the other a maximum (hard) limit. Soft limits may be changed by a process to any value that is less than or equal to the hard limit. A process may (irreversibly) lower its hard limit to any value that is greater than or equal to the soft limit. Only a process with an effective user ID of super-user can raise a hard limit. Both hard and soft limits can be changed in a single call to setrlimit() subject to the constraints described above. Limits may have an “infinite” value of <code>RLIM_INFINITY</code>. The <i>rlp</i> argument is a pointer to <code>struct rlimit</code> that includes the following members:</p> <pre>rlim_t rlim_cur; /* current (soft) limit */ rlim_t rlim_max; /* hard limit */</pre> <p>The type <code>rlim_t</code> is an arithmetic data type to which objects of type <code>int</code>, <code>size_t</code>, and <code>off_t</code> can be cast without loss of information.</p> <p>The possible resources, their descriptions, and the actions taken when the current limit is exceeded are summarized as follows:</p> <table border="0"> <tr> <td style="padding-right: 10px;"><code>RLIMIT_CORE</code></td> <td>The maximum size of a core file in bytes that may be created by a process. A limit of 0 will prevent the creation of a core file. The writing of a core file will terminate at this size.</td> </tr> <tr> <td style="padding-right: 10px;"><code>RLIMIT_CPU</code></td> <td>The maximum amount of CPU time in seconds used by a process. This is a soft limit only. The <code>SIGXCPU</code> signal is sent to the process. If the process is holding or ignoring <code>SIGXCPU</code>, the behavior is scheduling class defined.</td> </tr> <tr> <td style="padding-right: 10px;"><code>RLIMIT_DATA</code></td> <td>The maximum size of a process's heap in bytes. The <code>brk(2)</code> function will fail with <code>errno</code> set to <code>ENOMEM</code>.</td> </tr> </table>	<code>RLIMIT_CORE</code>	The maximum size of a core file in bytes that may be created by a process. A limit of 0 will prevent the creation of a core file. The writing of a core file will terminate at this size.	<code>RLIMIT_CPU</code>	The maximum amount of CPU time in seconds used by a process. This is a soft limit only. The <code>SIGXCPU</code> signal is sent to the process. If the process is holding or ignoring <code>SIGXCPU</code> , the behavior is scheduling class defined.	<code>RLIMIT_DATA</code>	The maximum size of a process's heap in bytes. The <code>brk(2)</code> function will fail with <code>errno</code> set to <code>ENOMEM</code> .
<code>RLIMIT_CORE</code>	The maximum size of a core file in bytes that may be created by a process. A limit of 0 will prevent the creation of a core file. The writing of a core file will terminate at this size.						
<code>RLIMIT_CPU</code>	The maximum amount of CPU time in seconds used by a process. This is a soft limit only. The <code>SIGXCPU</code> signal is sent to the process. If the process is holding or ignoring <code>SIGXCPU</code> , the behavior is scheduling class defined.						
<code>RLIMIT_DATA</code>	The maximum size of a process's heap in bytes. The <code>brk(2)</code> function will fail with <code>errno</code> set to <code>ENOMEM</code> .						

RLIMIT_FSIZE	<p>The maximum size of a file in bytes that may be created by a process. A limit of 0 will prevent the creation of a file. The <code>SIGXFSZ</code> signal is sent to the process. If the process is holding or ignoring <code>SIGXFSZ</code>, continued attempts to increase the size of a file beyond the limit will fail with <code>errno</code> set to <code>EFBIG</code>.</p>
RLIMIT_NOFILE	<p>One more than the maximum value that the system may assign to a newly created descriptor. This limit constrains the number of file descriptors that a process may create.</p>
RLIMIT_STACK	<p>The maximum size of a process's stack in bytes. The system will not automatically grow the stack beyond this limit.</p> <p>Within a process, <code>setrlimit()</code> will increase the limit on the size of your stack, but will not move current memory segments to allow for that growth. To guarantee that the process stack can grow to the limit, the limit must be altered prior to the execution of the process in which the new stack size is to be used.</p> <p>Within a multithreaded process, <code>setrlimit()</code> has no impact on the stack size limit for the calling thread if the calling thread is not the main thread. A call to <code>setrlimit()</code> for <code>RLIMIT_STACK</code> impacts only the main thread's stack, and should be made only from the main thread, if at all.</p> <p>The <code>SIGSEGV</code> signal is sent to the process. If the process is holding or ignoring <code>SIGSEGV</code>, or is catching <code>SIGSEGV</code> and has not made arrangements to use an alternate stack (see <code>sigaltstack(2)</code>), the disposition of <code>SIGSEGV</code> will be set to <code>SIG_DFL</code> before it is sent.</p>
RLIMIT_VMEM	<p>The maximum size of a process's mapped address space in bytes. If this limit is exceeded, the <code>brk(2)</code> and <code>mmap(2)</code> functions will fail with <code>errno</code> set to <code>ENOMEM</code>. In addition, the automatic stack growth will fail with the effects outlined above.</p>
RLIMIT_AS	<p>This is the maximum size of a process's total available memory, in bytes. If this limit is exceeded, the <code>brk(2)</code>, <code>malloc(3C)</code>, <code>mmap(2)</code> and <code>sbrk(2)</code> functions will fail with</p>

`errno` set to `ENOMEM`. In addition, the automatic stack growth will fail with the effects outlined above.

Because limit information is stored in the per-process information, the shell builtin `ulimit` command must directly execute this system call if it is to affect all future processes created by the shell.

The value of the current limit of the following resources affect these implementation defined parameters:

Limit	Implementation Defined Constant
<code>RLIMIT_FSIZE</code>	<code>FCHR_MAX</code>
<code>RLIMIT_NOFILE</code>	<code>OPEN_MAX</code>

When using the `getrlimit()` function, if a resource limit can be represented correctly in an object of type `rlim_t`, then its representation is returned; otherwise, if the value of the resource limit is equal to that of the corresponding saved hard limit, the value returned is `RLIM_SAVED_MAX`; otherwise the value returned is `RLIM_SAVED_CUR`.

When using the `setrlimit()` function, if the requested new limit is `RLIM_INFINITY`, the new limit will be "no limit"; otherwise if the requested new limit is `RLIM_SAVED_MAX`, the new limit will be the corresponding saved hard limit; otherwise, if the requested new limit is `RLIM_SAVED_CUR`, the new limit will be the corresponding saved soft limit; otherwise, the new limit will be the requested value. In addition, if the corresponding saved limit can be represented correctly in an object of type `rlim_t`, then it will be overwritten with the new limit.

The result of setting a limit to `RLIM_SAVED_MAX` or `RLIM_SAVED_CUR` is unspecified unless a previous call to `getrlimit()` returned that value as the soft or hard limit for the corresponding resource limit.

A limit whose value is greater than `RLIM_INFINITY` is permitted.

The `exec` family of functions also cause resource limits to be saved. See `exec(2)`.

RETURN VALUES

Upon successful completion, `getrlimit()` and `setrlimit()` return 0. Otherwise, these functions return -1 and set `errno` to indicate the error.

ERRORS

The `getrlimit()` and `setrlimit()` functions will fail if:

- EFAULT** The `rlp` argument points to an illegal address.
- EINVAL** An invalid *resource* was specified; or in a `setrlimit()` call, the new `rlim_cur` exceeds the new `rlim_max`.

EPERM The limit specified to **setrlimit()** would have raised the maximum limit value, and the effective user of the calling process is not super-user.

The **setrlimit()** function may fail if:

EINVAL The limit specified cannot be lowered because current usage is already higher than the limit.

USAGE The **getrlimit()** and **setrlimit()** functions have transitional interfaces for 64-bit file offsets. See **lf64(5)**.

SEE ALSO **brk(2)**, **exec(2)**, **fork(2)**, **open(2)**, **sigaltstack(2)**, **ulimit(2)**, **getdtablesize(3C)**, **malloc(3C)**, **signal(3C)**, **sysconf(3C)**, **lf64(5)**, **signal(5)**

NAME	getsid – get process group ID of session leader
SYNOPSIS	<pre>#include <unistd.h> pid_t getsid(pid_t pid);</pre>
DESCRIPTION	The getsid() function obtains the process group ID of the process that is the session leader of the process specified by <i>pid</i> . If <i>pid</i> is (pid_t) 0, it specifies the calling process.
RETURN VALUES	Upon successful completion, getsid() returns the process group ID of the session leader of the specified process. Otherwise, it returns (pid_t)-1 and sets <i>errno</i> to indicate the error.
ERRORS	The getsid() function will fail if: EPERM The process specified by <i>pid</i> is not in the same session as the calling process, and the implementation does not allow access to the process group ID of the session leader of that process from the calling process. ESRCH There is no process with a process ID equal to <i>pid</i> .
SEE ALSO	exec(2) , fork(2) , getpid(2) , getpgid(2) , setpgid(2) , setsid(2)

NAME	getuid, geteuid, getgid, getegid – get real user, effective user, real group, and effective group IDs				
SYNOPSIS	<pre>#include <sys/types.h> #include <unistd.h> uid_t getuid(void); uid_t geteuid(void); gid_t getgid(void); gid_t getegid(void);</pre>				
DESCRIPTION	<p>The getuid() function returns the real user ID of the calling process. The real user ID identifies the person who is logged in.</p> <p>The geteuid() function returns the effective user ID of the calling process. The effective user ID gives the process various permissions during execution of “set-user-ID” mode processes which use getuid() to determine the real user ID of the process that invoked them.</p> <p>The getgid() function returns the real group ID of the calling process.</p> <p>The getegid() function returns the effective group ID of the calling process.</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	intro(2) , setuid(2) , attributes(5)				

NAME	ioctl – control device
SYNOPSIS	<pre>#include <unistd.h> #include <stropts.h> int ioctl(int fildes, int request, /* arg */ ...);</pre>
DESCRIPTION	<p>The ioctl() function performs a variety of control functions on devices and STREAMS. For non-STREAMS files, the functions performed by this call are device-specific control functions. The <i>request</i> argument and an optional third argument with varying type are passed to the file designated by <i>fildes</i> and are interpreted by the device driver.</p> <p>For STREAMS files, specific functions are performed by the ioctl() function as described in streamio(7I).</p> <p>The <i>fildes</i> argument is an open file descriptor that refers to a device. The <i>request</i> argument selects the control function to be performed and depends on the device being addressed. The <i>arg</i> argument represents a third argument that has additional information that is needed by this specific device to perform the requested function. The data type of <i>arg</i> depends upon the particular control request, but it is either an <code>int</code> or a pointer to a device-specific data structure.</p> <p>In addition to device-specific and STREAMS functions, generic functions are provided by more than one device driver (for example, the general terminal interface.) See termio(7I).</p>
RETURN VALUES	Upon successful completion, the value returned depends upon the device control function, but must be a non-negative integer. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The ioctl() function will fail for any type of file if:</p> <p>EBADF The <i>fildes</i> argument is not a valid open file descriptor.</p> <p>EINTR A signal was caught during the execution of the ioctl() function.</p> <p>EINVAL The STREAM or multiplexer referenced by <i>fildes</i> is linked (directly or indirectly) downstream from a multiplexer.</p> <p>The ioctl() function will also fail if the device driver detects an error. In this case, the error is passed through ioctl() without change to the caller. A particular driver might not have all of the following error cases. Under the following conditions, requests to device drivers may fail and set <code>errno</code> to indicate the error</p> <p>EFAULT The <i>request</i> argument requires a data transfer to or from a buffer pointed to by <i>arg</i>, but <i>arg</i> points to an illegal address.</p>

EINVAL	The <i>request</i> or <i>arg</i> argument is not valid for this device.
EIO	Some physical I/O error has occurred.
ENOLINK	The <i>fildev</i> argument is on a remote machine and the link to that machine is no longer active.
ENOTTY	The <i>fildev</i> argument is not associated with a STREAMS device that accepts control functions.
ENXIO	The <i>request</i> and <i>arg</i> arguments are valid for this device driver, but the service requested can not be performed on this particular subdevice.
ENODEV	The <i>fildev</i> argument refers to a valid STREAMS device, but the corresponding device driver does not support the ioctl() function.

STREAMS errors are described in **streamio(7I)**.

SEE ALSO**streamio(7I)**, **termio(7I)**

NAME	kill – send a signal to a process or a group of processes
SYNOPSIS	<pre>#include <sys/types.h> #include <signal.h> int kill(pid_t pid, int sig);</pre>
DESCRIPTION	<p>The kill() function sends a signal 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 <i>pid</i>. The signal that is to be sent is specified by <i>sig</i> and is either one from the list given in <code>signal</code> (see signal(5)), or 0. If <i>sig</i> is 0 (the null signal), error checking is performed but no signal is actually sent. This can be used to check the validity of <i>pid</i>.</p> <p>The real or effective user ID of the sending process must match the real or saved (from one of functions in the <code>exec</code> family, see exec(2)) user ID of the receiving process unless the effective user ID of the sending process is super-user, (see intro(2)), or <i>sig</i> is <code>SIGCONT</code> and the sending process has the same session ID as the receiving process.</p> <p>If <i>pid</i> is greater than 0, <i>sig</i> will be sent to the process whose process ID is equal to <i>pid</i>.</p> <p>If <i>pid</i> is negative but not <code>(pid_t)-1</code>, <i>sig</i> will be sent to all processes whose process group ID is equal to the absolute value of <i>pid</i> and for which the process has permission to send a signal.</p> <p>If <i>pid</i> is 0, <i>sig</i> will be sent to all processes excluding special processes (see intro(2)) whose process group ID is equal to the process group ID of the sender.</p> <p>If <i>pid</i> is <code>(pid_t)-1</code> and the effective user ID of the sender is not super-user, <i>sig</i> will be sent to all processes excluding special processes whose real user ID is equal to the effective user ID of the sender.</p> <p>If <i>pid</i> is <code>(pid_t)-1</code> and the effective user ID of the sender is super-user, <i>sig</i> will be sent to all processes excluding special processes.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, no signal is sent, and <code>errno</code> is set to indicate the error.
ERRORS	<p>The kill() function will fail if:</p> <p>EINVAL The <i>sig</i> argument is not a valid signal number.</p> <p>EPERM The <i>sig</i> argument is <code>SIGKILL</code> and the <i>pid</i> argument is <code>(pid_t)1</code> (that is, the calling process does not have permission to send the signal to any of the processes specified by <i>pid</i>); or the effective user of the calling process</p>

does not match the real or saved user and is not super-user, and the calling process is not sending SIGCONT to a process that shares the same session ID.

ESRCH No process or process group can be found corresponding to that specified by *pid*.

USAGE

The `sigsend(2)` function provides a more versatile way to send signals to processes.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

`kill(1)`, `intro(2)`, `exec(2)`, `getpid(2)`, `getsid(2)`, `setpgrp(2)`, `sigaction(2)`, `sigsend(2)`, `signal(3C)`, `attributes(5)`, `signal(5)`

NAME	link – link to a file														
SYNOPSIS	<pre>#include <unistd.h> int link(const char *existing, const char *new);</pre>														
DESCRIPTION	<p>The link() function creates a new link (directory entry) for the existing file and increments its link count by one. The <i>existing</i> argument points to a path name naming an existing file. The <i>new</i> argument points to a pathname naming the new directory entry to be created.</p> <p>To create 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 <i>existing</i> must not be a directory.</p> <p>Upon successful completion, link() marks for update the <code>st_ctime</code> field of the file. Also, the <code>st_ctime</code> and <code>st_mtime</code> fields of the directory that contains the new entry are marked for update.</p>														
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, no link is created, and <code>errno</code> is set to indicate the error.														
ERRORS	<p>The link() function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EACCES</td> <td>A component of either path prefix denies search permission, or the requested link requires writing in a directory with a mode that denies write permission.</td> </tr> <tr> <td>EDQUOT</td> <td>The directory where the entry for the new link is being placed cannot be extended because the user's quota of disk blocks on that file system has been exhausted.</td> </tr> <tr> <td>EEXIST</td> <td>The link named by <i>new</i> exists.</td> </tr> <tr> <td>EFAULT</td> <td>The <i>existing</i> or <i>new</i> argument points to an illegal address.</td> </tr> <tr> <td>EINTR</td> <td>A signal was caught during the execution of the link() function.</td> </tr> <tr> <td>ELOOP</td> <td>Too many symbolic links were encountered in translating <i>path</i>.</td> </tr> <tr> <td>EMLINK</td> <td>The maximum number of links to a file would be exceeded.</td> </tr> </table>	EACCES	A component of either path prefix denies search permission, or the requested link requires writing in a directory with a mode that denies write permission.	EDQUOT	The directory where the entry for the new link is being placed cannot be extended because the user's quota of disk blocks on that file system has been exhausted.	EEXIST	The link named by <i>new</i> exists.	EFAULT	The <i>existing</i> or <i>new</i> argument points to an illegal address.	EINTR	A signal was caught during the execution of the link() function.	ELOOP	Too many symbolic links were encountered in translating <i>path</i> .	EMLINK	The maximum number of links to a file would be exceeded.
EACCES	A component of either path prefix denies search permission, or the requested link requires writing in a directory with a mode that denies write permission.														
EDQUOT	The directory where the entry for the new link is being placed cannot be extended because the user's quota of disk blocks on that file system has been exhausted.														
EEXIST	The link named by <i>new</i> exists.														
EFAULT	The <i>existing</i> or <i>new</i> argument points to an illegal address.														
EINTR	A signal was caught during the execution of the link() function.														
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .														
EMLINK	The maximum number of links to a file would be exceeded.														

ENAMETOOLONG	The length of the <i>existing</i> or <i>new</i> argument exceeds <code>PATH_MAX</code> , or the length of a <i>existing</i> or <i>new</i> component exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.
ENOENT	The <i>existing</i> or <i>new</i> argument is a null pathname; a component of either path prefix does not exist; or the file named by <i>existing</i> does not exist.
ENOLINK	The <i>existing</i> or <i>new</i> argument points to a remote machine and the link to that machine is no longer active.
ENOSPC	The directory that would contain the link cannot be extended.
ENOTDIR	A component of either path prefix is not a directory.
EPERM	The file named by <i>existing</i> is a directory and the effective user of the calling process is not super-user.
EROFS	The requested link requires writing in a directory on a read-only file system.
EXDEV	The link named by <i>new</i> and the file named by <i>existing</i> are on different logical devices (file systems).

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

`symlink(2)`, `unlink(2)`, `attributes(5)`

NAME	llseek – move extended read/write file pointer
SYNOPSIS	<pre>#include <sys/types.h> #include <unistd.h> offset_t llseek(int <i>fildev</i>, offset_t <i>offset</i>, int <i>whence</i>);</pre>
DESCRIPTION	<p>The llseek() function sets the 64-bit extended file pointer associated with the open file descriptor specified by <i>fildev</i> as follows:</p> <ul style="list-style-type: none"> ■ If <i>whence</i> is <code>SEEK_SET</code>, the pointer is set to <i>offset</i> bytes. ■ If <i>whence</i> is <code>SEEK_CUR</code>, the pointer is set to its current location plus <i>offset</i>. ■ If <i>whence</i> is <code>SEEK_END</code>, the pointer is set to the size of the file plus <i>offset</i>. <p>Although each file has a 64-bit file pointer associated with it, existing file system types do not support the full range of 64-bit offsets. In particular, non-device files remain limited to offsets of less than two gigabytes. Device drivers may support offsets of up to 1024 gigabytes for device special files.</p> <p>Some devices are incapable of seeking. The value of the file pointer associated with such a device is undefined.</p>
RETURN VALUES	<p>Upon successful completion, llseek() returns the resulting pointer location as measured in bytes from the beginning of the file. Remote file descriptors are the only ones that allow negative file pointers. Otherwise, <code>-1</code> is returned, the file pointer remains unchanged, and <code>errno</code> is set to indicate the error.</p>
ERRORS	<p>The llseek() function will fail if:</p> <p>EBADF The <i>fildev</i> argument is not an open file descriptor.</p> <p>EINVAL The <i>whence</i> argument is not <code>SEEK_SET</code>, <code>SEEK_CUR</code>, or <code>SEEK_END</code>; the <i>offset</i> argument is not a valid offset for this file system type; or the <i>fildev</i> argument is not a remote file descriptor, and the resulting file pointer would be negative.</p> <p>ESPIPE The <i>fildev</i> argument is associated with a pipe or FIFO.</p>
SEE ALSO	<code>creat(2)</code> , <code>dup(2)</code> , <code>fcntl(2)</code> , <code>lseek(2)</code> , <code>open(2)</code>

NAME	lseek – move read/write file pointer								
SYNOPSIS	<pre>#include <sys/types.h> #include <unistd.h> off_t lseek(int <i>fd</i>, off_t <i>offset</i>, int <i>whence</i>);</pre>								
DESCRIPTION	<p>The lseek() function sets the file pointer associated with the open file descriptor specified by <i>fd</i> as follows:</p> <ul style="list-style-type: none"> ■ If <i>whence</i> is <code>SEEK_SET</code>, the pointer is set to <i>offset</i> bytes. ■ If <i>whence</i> is <code>SEEK_CUR</code>, the pointer is set to its current location plus <i>offset</i>. ■ If <i>whence</i> is <code>SEEK_END</code>, the pointer is set to the size of the file plus <i>offset</i>. <p>The symbolic constants <code>SEEK_SET</code>, <code>SEEK_CUR</code>, and <code>SEEK_END</code> are defined in the header <code><unistd.h></code>.</p> <p>Some devices are incapable of seeking. The value of the file pointer associated with such a device is undefined.</p> <p>The lseek() function allows the file pointer to be set beyond the existing data in the file. If data are later written at this point, subsequent reads in the gap between the previous end of data and the newly written data will return bytes of value 0 until data are written into the gap.</p> <p>If <i>fd</i> is a remote file descriptor and <i>offset</i> is negative, lseek() returns the file pointer even if it is negative. The lseek() function will not, by itself, extend the size of a file.</p>								
RETURN VALUES	<p>Upon successful completion, the resulting offset, as measured in bytes from the beginning of the file, is returned. Otherwise, <code>(off_t)-1</code> is returned, the file offset remains unchanged, and <code>errno</code> is set to indicate the error.</p>								
ERRORS	<p>The lseek() function will fail if:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">EBADF</td> <td>The <i>fd</i> argument is not an open file descriptor.</td> </tr> <tr> <td>EINVAL</td> <td>The <i>whence</i> argument is not <code>SEEK_SET</code>, <code>SEEK_CUR</code>, or <code>SEEK_END</code>; or the <i>fd</i> argument is not a remote file descriptor and the resulting file pointer would be negative.</td> </tr> <tr> <td>E_OVERFLOW</td> <td>The resulting file offset would be a value which cannot be represented correctly in an object of type <code>off_t</code> for regular files.</td> </tr> <tr> <td>ESPIPE</td> <td>The <i>fd</i> argument is associated with a pipe, a FIFO, or a socket.</td> </tr> </table>	EBADF	The <i>fd</i> argument is not an open file descriptor.	EINVAL	The <i>whence</i> argument is not <code>SEEK_SET</code> , <code>SEEK_CUR</code> , or <code>SEEK_END</code> ; or the <i>fd</i> argument is not a remote file descriptor and the resulting file pointer would be negative.	E_OVERFLOW	The resulting file offset would be a value which cannot be represented correctly in an object of type <code>off_t</code> for regular files.	ESPIPE	The <i>fd</i> argument is associated with a pipe, a FIFO, or a socket.
EBADF	The <i>fd</i> argument is not an open file descriptor.								
EINVAL	The <i>whence</i> argument is not <code>SEEK_SET</code> , <code>SEEK_CUR</code> , or <code>SEEK_END</code> ; or the <i>fd</i> argument is not a remote file descriptor and the resulting file pointer would be negative.								
E_OVERFLOW	The resulting file offset would be a value which cannot be represented correctly in an object of type <code>off_t</code> for regular files.								
ESPIPE	The <i>fd</i> argument is associated with a pipe, a FIFO, or a socket.								

USAGE The **lseek()** function has a transitional interface for 64-bit file offsets. See **lf64(5)**.

In multithreaded applications, using **lseek()** in conjunction with a **read(2)** or **write(2)** call on a file descriptor shared by more than one thread is not an atomic operation. To ensure atomicity, use **pread()** or **pwrite()**.

ATTRIBUTES See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO **creat(2)**, **dup(2)**, **fcntl(2)**, **open(2)**, **read(2)**, **write(2)**, **attributes(5)**, **lf64(5)**

NAME	<code>_lwp_cond_signal</code> , <code>_lwp_cond_broadcast</code> – signal a condition variable
SYNOPSIS	<pre>#include <sys/lwp.h> int _lwp_cond_signal(lwp_cond_t * cvp); int _lwp_cond_broadcast(lwp_cond_t * cvp);</pre>
DESCRIPTION	<p>The <code>_lwp_cond_signal()</code> function unblocks one LWP that is blocked on the LWP condition variable pointed to by <code>cvp</code> .</p> <p>The <code>_lwp_cond_broadcast()</code> function unblocks all LWPs that are blocked on the LWP condition variable pointed to by <code>cvp</code> .</p> <p>If no LWPs are blocked on the LWP condition variable, then <code>_lwp_cond_signal()</code> and <code>_lwp_cond_broadcast()</code> have no effect.</p> <p>Both functions should be called under the protection of the same LWP mutex lock that is used with the LWP condition variable being signaled. Otherwise, the condition variable may be signalled between the test of the associated condition and blocking in <code>_lwp_cond_wait()</code> . This can cause an infinite wait.</p>
RETURN VALUES	Upon successful completion, 0 is returned. A non-zero value indicates an error.
ERRORS	The <code>_lwp_cond_signal()</code> and <code>_lwp_cond_broadcast()</code> functions will fail if: EINVAL The <code>cvp</code> argument points to an invalid LWP condition variable. EFAULT The <code>cvp</code> argument points to an invalid address.
SEE ALSO	<code>_lwp_cond_wait(2)</code> , <code>_lwp_mutex_lock(2)</code>

NAME	_lwp_cond_wait, _lwp_cond_timedwait – wait on a condition variable
SYNOPSIS	<pre>#include <sys/lwp.h> int _lwp_cond_wait(lwp_cond_t * cvp, lwp_mutex_t * mp); int _lwp_cond_timedwait(lwp_cond_t * cvp, lwp_mutex_t * mp, timestruc_t * abstime);</pre>
DESCRIPTION	<p>These functions are used to wait for the occurrence of a condition represented by an LWP condition variable. LWP condition variables must be initialized to 0 before use.</p> <p>The _lwp_cond_wait() function atomically releases the LWP mutex pointed to by <i>mp</i> and causes the calling LWP to block on the LWP condition variable pointed to by <i>cvp</i>. The blocked LWP may be awakened by _lwp_cond_signal(2), _lwp_cond_broadcast(2), or when interrupted by delivery of a signal. Any change in value of a condition associated with the condition variable cannot be inferred by the return of _lwp_cond_wait() and any such condition must be re-evaluated.</p> <p>The _lwp_cond_timedwait() function is similar to _lwp_cond_wait(), except that the calling LWP will not block past the time of day specified by <i>abstime</i>. If the time of day becomes greater than <i>abstime</i>, _lwp_cond_timedwait() returns with the error code ETIME.</p> <p>The _lwp_cond_wait() and _lwp_cond_timedwait() functions always return with the mutex locked and owned by the calling lightweight process.</p>
RETURN VALUES	Upon successful completion, 0 is returned. A non-zero value indicates an error.
ERRORS	<p>If any of the following conditions are detected, _lwp_cond_wait() and _lwp_cond_timedwait() fail and return the corresponding value:</p> <p>EINVAL The <i>cvp</i> argument points to an invalid LWP condition variable or the <i>mp</i> argument points to an invalid LWP mutex.</p> <p>EFAULT The <i>mp</i>, <i>cvp</i>, or <i>abstime</i> argument points to an illegal address.</p> <p>If any of the following conditions occur, _lwp_cond_wait() and _lwp_cond_timedwait() fail and return the corresponding value:</p> <p>EINTR The call was interrupted by a signal or fork(2).</p> <p>If any of the following conditions occur, _lwp_cond_timedwait() fails and returns the corresponding value:</p> <p>ETIME The time specified in <i>abstime</i> has passed.</p>

EXAMPLES

EXAMPLE 1 Using the `_lwp_cond_wait()` function in a loop testing some condition.

The `_lwp_cond_wait()` function is normally used in a loop testing some condition, as follows:

```
lwp_mutex_t m;
lwp_cond_t cv;
int cond;
(void) _lwp_mutex_lock(&m);
while (cond == FALSE) {

    (void) _lwp_cond_wait(&cv, &m);
}
(void) _lwp_mutex_unlock(&m);
```

EXAMPLE 2 Using the `_lwp_cond_timedwait()` function in a loop testing some condition.

The `_lwp_cond_timedwait()` function is also normally used in a loop testing some condition. It uses an absolute timeout value as follows:

```
timestruc_t to;
lwp_mutex_t m;
lwp_cond_t cv;
int cond, err;
(void) _lwp_mutex_lock(&m);
to.tv_sec = time(NULL) + TIMEOUT;
to.tv_nsec = 0;
while (cond == FALSE) {
    err = _lwp_cond_timedwait(&cv, &m, &to);
    if (err == ETIME) {
        /* timeout, do something */
        break;
        SENDwhom}
}
(void) _lwp_mutex_unlock(&m);
```

This example sets a bound on the total wait time even though the `_lwp_cond_timedwait()` may return several times due to the condition being signalled or the wait being interrupted.

SEE ALSO

`_lwp_cond_broadcast(2)`, `_lwp_cond_signal(2)`, `_lwp_kill(2)`,
`_lwp_mutex_lock(2)`, `fork(2)`, `kill(2)`

NAME	_lwp_create – create a new light-weight process
SYNOPSIS	<pre>#include <sys/lwp.h> int _lwp_create(ucontext_t *contextp, unsigned long flags, lwpid_t *new_lwp);</pre>
DESCRIPTION	<p>The _lwp_create() function adds a lightweight process (LWP) to the current process. The <i>contextp</i> argument specifies the initial signal mask, stack, and machine context (including the program counter and stack pointer) for the new LWP. The new LWP inherits the scheduling class and priority of the caller.</p> <p>If _lwp_create() is successful and <i>new_lwp</i> is not null, the ID of the new LWP is stored in the location pointed to by <i>new_lwp</i>.</p> <p>The <i>flags</i> argument specifies additional attributes for the new LWP. The value in <i>flags</i> is constructed by the bitwise inclusive OR operation of the following values:</p> <p>LWP_DETACHED The LWP is created detached.</p> <p>LWP_SUSPENDED The LWP is created suspended.</p> <p>__LWP_ASLWP The LWP created is the ASLWP (Asynchronous Signals LWP) (see signal(5)). The ASLWP should always be created with all signals blocked. If __LWP_ASLWP is specified, then the LWP created is the special, designated LWP that handles signals sent to a multithreaded process (ASLWP). There can be only one ASLWP in a multithreaded process, so the creation of another ASLWP will return EINVAL. It should never exit by way of _lwp_exit(2) or exit(2). This is a reserved flag and should not be used by any user program. It is documented here for the sake of completion and not for use by an application.</p> <p>If LWP_DETACHED is specified, then the LWP is created in the <i>detached</i> state. Otherwise the LWP is created in the undetached state. The ID (and system resources) associated with a detached LWP can be automatically reclaimed when the LWP exits. The ID of an undetached LWP cannot be reclaimed until it exits and another LWP has reported its termination by way of _lwp_wait(2). This allows the waiting LWP to determine that the waited for LWP has terminated and to reclaim any process resources that it was using.</p> <p>If LWP_SUSPENDED is specified, then the LWP is created in a suspended state. This allows the creator to change the LWP's inherited attributes before it starts to execute. The suspended LWP can only be resumed by way of _lwp_continue(2). If LWP_SUSPENDED is not specified the LWP can begin to run immediately after it has been created.</p>
RETURN VALUES	Upon successful completion, 0 is returned. A non-zero value indicates an error.

ERRORS

If any of the following conditions are detected, **_lwp_create()** fails and returns the corresponding value:

- EFAULT** Either the *context* parameter or the *new_lwp* parameter point to invalid addresses.
- EAGAIN** A system limit is exceeded, (for example, too many LWP were created for this real user ID).
- EINVAL** The `__LWP_AS_LWP` flag was used to create more than one ASLWP in the process. There can be only one ASLWP within a process.

EXAMPLES

EXAMPLE 1 How a stack is allocated to a new LWP.

This example shows how a stack is allocated to a new LWP. The **_lwp_makecontext()** function is used to set up the *context* parameter so that the new LWP begins executing a function.

```

contextp = (ucontext_t *)malloc(sizeof(ucontext_t));
stackbase = malloc(stacksize);
sigprocmask(SIGSETMASK, NULL, &contextp->uc_sigmask);
_lwp_makecontext(contextp, func, arg, private, stackbase, stacksize);
error = _lwp_create(contextp, NULL, &new_lwp);

```

SEE ALSO

_lwp_cond_timedwait(2), _lwp_continue(2), _lwp_exit(2), _lwp_makecontext(2), _lwp_wait(2), alarm(2), exit(2), poll(2), sleep(3C), thr_create(3T), signal(5), ucontext(5)

NOTES

Applications should use bound threads rather than the `_lwp_*()` functions (see **thr_create(3T)**). Using LWPs directly is not advised because libraries are only safe to use with threads, not LWPs.

In Solaris releases 2.5 through 7, the signal `SIGALRM` is defined to be per-process. This does not affect the behavior of single-threaded or multithreaded applications. If the application was using LWPs directly, and was relying on **alarm(2)** or **sleep(3C)**, then the application's behavior might be impacted. The calling LWP will not necessarily be the recipient of the `SIGALRM` signal when `SIGALRM` is sent to the process. You might have to use a substitute like **poll(2)**, or **_lwp_cond_timedwait(2)** to simulate the old per-LWP semantic of `SIGALRM`.

NAME	<code>_lwp_exit</code> - terminate the calling LWP
SYNOPSIS	<pre>#include <sys/lwp.h> void _lwp_exit(void);</pre>
DESCRIPTION	<p>The <code>_lwp_exit()</code> function causes the calling LWP to terminate. If it is the last LWP in the process, then the process exits with a status of 0 (see <code>exit(2)</code>).</p> <p>If the LWP was created undetached, it is transformed into a “zombie LWP” that retains at least the LWP’s ID until it is waited for (see <code>_lwp_wait(2)</code>). Otherwise, its ID and system resources may be reclaimed immediately.</p>
SEE ALSO	<code>_lwp_create(2)</code> , <code>_lwp_wait(2)</code> , <code>exit(2)</code>

NAME	<code>_lwp_info</code> - return the time-accounting information of a single LWP
SYNOPSIS	<pre>#include <sys/time.h> #include <sys/lwp.h> int _lwp_info(struct lwpinfo *buffer);</pre>
DESCRIPTION	<p>The <code>_lwp_info()</code> function fills the <code>lwpinfo</code> structure pointed to by <i>buffer</i> with time-accounting information pertaining to the calling LWP. This call may be extended in the future to return other information to the <code>lwpinfo</code> structure as needed. The <code>lwpinfo</code> structure in <code><sys/lwp.h></code> includes the following members:</p> <pre>timestruct_t lwp_utime; timestruct_t lwp_stime;</pre> <p>The <code>lwp_utime</code> member is the CPU time used while executing instructions in the user space of the calling LWP.</p> <p>The <code>lwp_stime</code> member is the CPU time used by the system on behalf of the calling LWP.</p>
RETURN VALUES	Upon successful completion, <code>_lwp_info()</code> returns 0 and fills in the <code>lwpinfo</code> structure pointed to by <i>buffer</i> .
ERRORS	If the following condition is detected, <code>_lwp_info()</code> returns the corresponding value: EFAULT The <i>buffer</i> argument points to an illegal address. Additionally, the <code>_lwp_info()</code> function will fail for 32-bit interfaces if: Eoverflow The size of the <code>tv_sec</code> member of the <code>timestruct_t</code> type pointed to by <code>lwp_utime</code> and <code>lwp_stime</code> is too small to contain the correct number of seconds.
SEE ALSO	<code>times(2)</code>

NAME	<code>_lwp_kill</code> – send a signal to a LWP
SYNOPSIS	<pre>#include <sys/lwp.h> #include <signal.h> int _lwp_kill(lwpid_t target_lwp, int sig);</pre>
DESCRIPTION	<p>The <code>_lwp_kill()</code> function sends a signal to the LWP specified by <i>target_lwp</i>. The signal that is to be sent is specified by <i>sig</i> and must be one from the list given in <code>signal(5)</code>. If <i>sig</i> is 0 (the null signal), error checking is performed but no signal is actually sent. This can be used to check the validity of <i>target_lwp</i>.</p> <p>The <i>target_lwp</i> must be an LWP within the same process as the calling LWP.</p>
RETURN VALUES	Upon successful completion, 0 is returned. A non-zero value indicates an error.
ERRORS	If any of the following conditions occur, <code>_lwp_kill()</code> fails and returns the corresponding value:
EINVAL	The <i>sig</i> argument is not a valid signal number.
ESRCH	The <i>target_lwp</i> argument cannot be found in the current process.
SEE ALSO	<code>kill(2)</code> , <code>sigaction(2)</code> , <code>sigprocmask(2)</code> , <code>signal(5)</code>

NAME	<code>_lwp_makecontext</code> - initialize an LWP context
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/lwp.h> #include <ucontext.h> void _lwp_makecontext(ucontext_t *ucp, void (*start_routine)(void *), void *arg, void *private, caddr_t stack_base, size_t stack_size);</pre>
DESCRIPTION	<p>The <code>_lwp_makecontext()</code> function initializes the user context structure pointed to by <code>ucp</code>. The user context is defined by <code>ucontext(5)</code>. The resulting user context can be used by <code>_lwp_create(2)</code> for specifying the initial state of the new LWP. The user context is set up to start executing the function <code>start_routine</code> with a single argument, <code>arg</code>, and to call <code>_lwp_exit(2)</code> if <code>start_routine</code> returns. The new LWP will use the storage starting at <code>stack_base</code> and continuing for <code>stack_size</code> bytes as an execution stack. The initial value in LWP-private memory will be set to <code>private</code> (see <code>_lwp_setprivate(2)</code>). The signal mask in the user context is not initialized.</p>
SEE ALSO	<code>_lwp_create(2)</code> , <code>_lwp_exit(2)</code> , <code>_lwp_setprivate(2)</code> , <code>ucontext(5)</code>

NAME	_lwp_mutex_lock, _lwp_mutex_unlock, _lwp_mutex_trylock – mutual exclusion
SYNOPSIS	<pre>#include <sys/lwp.h> int _lwp_mutex_lock(lwp_mutex_t * mp); int _lwp_mutex_trylock(lwp_mutex_t * mp); int _lwp_mutex_unlock(lwp_mutex_t * mp);</pre>
DESCRIPTION	<p>These functions serialize the execution of lightweight processes. They are useful for ensuring that only one lightweight process can execute a critical section of code at any one time (mutual exclusion). LWP mutexes must be initialized to 0 before use.</p> <p>The _lwp_mutex_lock() function locks the LWP mutex pointed to by <i>mp</i> . If the mutex is already locked, the calling LWP blocks until the mutex becomes available. When _lwp_mutex_lock() returns, the mutex is locked and the calling LWP is the "owner".</p> <p>The _lwp_mutex_trylock() function attempts to lock the mutex. If the mutex is already locked it returns with an error. If the mutex is unlocked, it is locked and _lwp_mutex_trylock() returns.</p> <p>The _lwp_mutex_unlock() function unlocks a locked mutex. The mutex must be locked and the calling LWP must be the one that last locked the mutex (the owner). If any other LWPs are waiting for the mutex to become available, one of them is unblocked.</p>
RETURN VALUES	Upon successful completion, 0 is returned. A non-zero value indicates an error.
ERRORS	<p>If any of the following conditions are detected, _lwp_mutex_lock() , _lwp_mutex_trylock() , and _lwp_mutex_unlock() fail and return the corresponding value:</p> <p>EINVAL The <i>mp</i> argument points to an invalid LWP mutex.</p> <p>EFAULT The <i>mp</i> argument points to an illegal address.</p> <p>If any of the following conditions occur, _lwp_mutex_trylock() fails and returns the corresponding value:</p> <p>EBUSY The <i>mp</i> argument points to a locked mutex.</p>
SEE ALSO	<code>intro(2)</code> , <code>_lwp_cond_wait(2)</code>

NAME	<code>_lwp_self</code> - get LWP identifier
SYNOPSIS	<code>#include <sys/lwp.h></code> <code>lwpid_t _lwp_self(void);</code>
DESCRIPTION	The <code>_lwp_self()</code> function returns the ID of the calling LWP.
SEE ALSO	<code>_lwp_create(2)</code>

NAME	<code>_lwp_sema_wait</code> , <code>_lwp_sema_trywait</code> , <code>_lwp_sema_init</code> , <code>_lwp_sema_post</code> – semaphore operations
SYNOPSIS	<pre>#include <sys/lwp.h> int _lwp_sema_wait(lwp_sema_t * sema); int _lwp_sema_trywait(lwp_sema_t * sema); int _lwp_sema_init(lwp_sema_t * sema, int count); int _lwp_sema_post(lwp_sema_t * sema);</pre>
DESCRIPTION	<p>Conceptually, a semaphore is a non-negative integer count that is atomically incremented and decremented. Typically this represents the number of resources available. The <code>_lwp_sema_init()</code> function initializes the count, <code>_lwp_sema_post()</code> atomically increments the count, and <code>_lwp_sema_wait()</code> waits for the count to become greater than 0 and then atomically decrements it.</p> <p>LWP semaphores must be initialized before use. The <code>_lwp_sema_init()</code> function initializes the count associated with the LWP semaphore pointed to by <code>sema</code> to <code>count</code>.</p> <p>The <code>_lwp_sema_wait()</code> function blocks the calling LWP until the semaphore count becomes greater than 0 and then atomically decrements it.</p> <p>The <code>_lwp_sema_trywait()</code> function atomically decrements the count if it is greater than zero. Otherwise it returns an error.</p> <p>The <code>_lwp_sema_post()</code> function atomically increments the semaphore count. If there are any LWPs blocked on the semaphore, one is unblocked.</p>
RETURN VALUES	Upon successful completion, 0 is returned. A non-zero value indicates an error.
ERRORS	<p>The <code>_lwp_sema_init()</code>, <code>_lwp_sema_trywait()</code>, <code>_lwp_sema_wait()</code>, and <code>_lwp_sema_post()</code> functions will fail if:</p> <p>EINVAL The <code>sema</code> argument points to an invalid semaphore.</p> <p>EFAULT The <code>sema</code> argument points to an illegal address.</p> <p>The <code>_lwp_sema_wait()</code> function will fail if:</p> <p>EINTR The function execution was interrupted by a signal or <code>fork(2)</code>.</p> <p>The <code>_lwp_sema_trywait()</code> function will fail if:</p> <p>EBUSY The function was called on a semaphore with a zero count.</p> <p>The <code>_lwp_sema_post()</code> function will fail if:</p> <p>EOVERFLOW The value of the <code>sema</code> argument exceeds <code>SEM_VALUE_MAX</code>.</p>

`_lwp_sema_wait(2)`

System Calls

SEE ALSO

`fork(2)`

NAME	<code>_lwp_setprivate</code> , <code>_lwp_getprivate</code> – set or get LWP specific storage
SYNOPSIS	<pre>#include <sys/lwp.h> void _lwp_setprivate(void * <i>buffer</i>); void * _lwp_getprivate(void);</pre>
DESCRIPTION	<p>The <code>_lwp_setprivate()</code> function stores the value specified by <i>buffer</i> in LWP-private memory that is unique to the calling LWP. This is typically used by thread library implementations to maintain a pointer to information about the thread currently running on the calling LWP.</p> <p>The <code>_lwp_getprivate()</code> function returns the value stored in LWP-private memory.</p>
SEE ALSO	<code>_lwp_makecontext(2)</code>

NAME	<code>_lwp_suspend</code> , <code>_lwp_continue</code> – continue or suspend LWP execution
SYNOPSIS	<pre>#include <sys/lwp.h> int _lwp_suspend(lwpid_t target_lwp); int _lwp_continue(lwpid_t target_lwp);</pre>
DESCRIPTION	<p>The <code>_lwp_suspend()</code> function immediately suspends the execution of the LWP specified by <code>target_lwp</code>. On successful return from <code>_lwp_suspend()</code>, <code>target_lwp</code> is no longer executing. Once a thread is suspended, subsequent calls to <code>_lwp_suspend()</code> have no affect.</p> <p>The <code>_lwp_continue()</code> function resumes the execution of a suspended LWP. Once a suspended LWP is continued, subsequent calls to <code>_lwp_continue()</code> have no effect.</p> <p>A suspended LWP will not be awakened by a signal. The signal stays pending until the execution of the LWP is resumed by <code>_lwp_continue()</code>.</p>
RETURN VALUES	Upon successful completion, 0 is returned. A non-zero value indicates an error.
ERRORS	<p>If the following condition occurs, <code>_lwp_suspend()</code> and <code>_lwp_continue()</code> fail and return the corresponding value:</p> <p>ESRCH The <code>target_lwpid</code> argument cannot be found in the current process.</p> <p>If the following condition is detected, <code>_lwp_suspend()</code> fails and returns the corresponding value:</p> <p>EDEADLK Suspending <code>target_lwpid</code> will cause all LWPs in the process to be suspended.</p>
SEE ALSO	<code>_lwp_create(2)</code>

NAME	<code>_lwp_wait</code> – wait for a LWP to terminate
SYNOPSIS	<pre>#include <sys/lwp.h> int _lwp_wait(lwpid_t wait_for, lwpid_t *departed_lwp);</pre>
DESCRIPTION	<p>The <code>_lwp_wait()</code> function blocks the current LWP until the LWP specified by <code>wait_for</code> terminates. If the specified LWP terminated prior to the call to <code>_lwp_wait()</code>, then <code>_lwp_wait()</code> returns immediately. If <code>wait_for</code> is <code>NULL</code>, then <code>_lwp_wait()</code> waits for any undetached LWP in the current process. If <code>wait_for</code> is not <code>NULL</code>, then it must specify an undetached LWP in the current process. If <code>departed_lwp</code> is not <code>NULL</code>, then it points to location where the ID of the exited LWP is stored (see <code>_lwp_exit(2)</code>).</p> <p>When an LWP exits and there are one or more LWPs in this process waiting for this specific LWP to exit, then one of the waiting LWPs is unblocked and it returns from <code>_lwp_wait()</code> successfully. Any other LWPs waiting for this same LWP to exit are also unblocked, however, they return from <code>_lwp_wait()</code> with an error (<code>ESRCH</code>) indicating the waited for LWP no longer exists. If there are no LWPs in this process waiting for this specific LWP to exit but there are one or more LWPs waiting for any LWP to exit, then one of the waiting LWPs is unblocked and it returns from <code>_lwp_wait()</code> successfully.</p> <p>The ID of an LWP that has exited may be reused via <code>_lwp_create()</code> after the LWP has been successfully waited for.</p>
RETURN VALUES	Upon successful completion, 0 is returned. A non-zero value indicates an error.
ERRORS	<p>If any of the following conditions are detected, <code>_lwp_wait()</code> fails and returns the corresponding value:</p> <p>EINTR The <code>_lwp_wait()</code> function was interrupted by a signal.</p> <p>EDEADLK All LWPs in this process would be blocked waiting for LWPs to terminate, or the calling LWP is attempting to wait for itself.</p> <p>If any of the following conditions occur, <code>_lwp_wait()</code> fails and returns the corresponding value:</p> <p>ESRCH The <code>wait_for</code> argument cannot be found in the current process or it was detached.</p>
SEE ALSO	<code>_lwp_create(2)</code> , <code>_lwp_exit(2)</code>

NAME	mемсnt1 – memory management control
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/mman.h> int мемсnt1(caddr_t addr, size_t len, int cmd, caddr_t arg, int attr, int mask);</pre>
DESCRIPTION	<p>The mемсnt1() function allows the calling process to apply a variety of control operations over the address space identified by the mappings established for the address range [<i>addr</i>, <i>addr + len</i>).</p> <p>The <i>addr</i> argument must be a multiple of the pagesize as returned by sysconf(3C). The scope of the control operations can be further defined with additional selection criteria (in the form of attributes) according to the bit pattern contained in <i>attr</i>.</p> <p>The following attributes specify page mapping selection criteria:</p> <p>SHARED Page is mapped shared.</p> <p>PRIVATE Page is mapped private.</p> <p>The following attributes specify page protection selection criteria. The selection criteria are constructed by a bitwise OR operation on the attribute bits and must match exactly.</p> <p>PROT_READ Page can be read.</p> <p>PROT_WRITE Page can be written.</p> <p>PROT_EXEC Page can be executed.</p> <p>The following criteria may also be specified:</p> <p>PROC_TEXT Process text.</p> <p>PROC_DATA Process data.</p> <p>The PROC_TEXT attribute specifies all privately mapped segments with read and execute permission, and the PROC_DATA attribute specifies all privately mapped segments with write permission.</p> <p>Selection criteria can be used to describe various abstract memory objects within the address space on which to operate. If an operation shall not be constrained by the selection criteria, <i>attr</i> must have the value 0.</p> <p>The operation to be performed is identified by the argument <i>cmd</i>. The symbolic names for the operations are defined in <code><sys/mman.h></code> as follows:</p> <p>MC_LOCK Lock in memory all pages in the range with attributes <i>attr</i>. A given page may be locked multiple times through different mappings; however, within a given mapping, page locks do not nest. Multiple lock operations on the same address in the same process will all be removed with a single unlock</p>

operation. A page locked in one process and mapped in another (or visible through a different mapping in the locking process) is locked in memory as long as the locking process does neither an implicit nor explicit unlock operation. If a locked mapping is removed, or a page is deleted through file removal or truncation, an unlock operation is implicitly performed. If a writable `MAP_PRIVATE` page in the address range is changed, the lock will be transferred to the private page.

At present *arg* is unused, but must be 0 to ensure compatibility with potential future enhancements.

`MC_LOCKAS` Lock in memory all pages mapped by the address space with attributes *attr*. At present *addr* and *len* are unused, but must be `NULL` and 0 respectively, to ensure compatibility with potential future enhancements. The *arg* argument is a bit pattern built from the flags:

`MCL_CURRENT` Lock current mappings

`MCL_FUTURE` Lock future mappings

The value of *arg* determines whether the pages to be locked are those currently mapped by the address space, those that will be mapped in the future, or both. If `MCL_FUTURE` is specified, then all mappings subsequently added to the address space will be locked, provided sufficient memory is available.

`MC_SYNC` Write to their backing storage locations all modified pages in the range with attributes *attr*. Optionally, invalidate cache copies. The backing storage for a modified `MAP_SHARED` mapping is the file the page is mapped to; the backing storage for a modified `MAP_PRIVATE` mapping is its swap area. The *arg* argument is a bit pattern built from the flags used to control the behavior of the operation:

`MS_ASYNC` perform asynchronous writes

`MS_SYNC` perform synchronous writes

`MS_INVALIDATE` invalidate mappings

`MS_ASYNC` returns immediately once all write operations are scheduled; with `MS_SYNC` the function will not return until all write operations are completed.

MS_INVALIDATE invalidates all cached copies of data in memory, so that further references to the pages will be obtained by the system from their backing storage locations. This operation should be used by applications that require a memory object to be in a known state.

MC_UNLOCK Unlock all pages in the range with attributes *attr*. At present *arg* is unused, but must be 0 to ensure compatibility with potential future enhancements.

MC_UNLOCKAS Remove address space memory locks, and locks on all pages in the address space with attributes *attr*. At present *addr*, *len*, and *arg* are unused, but must be NULL, 0 and 0 respectively, to ensure compatibility with potential future enhancements.

The *mask* argument must be 0; it is reserved for future use.

Locks established with the lock operations are not inherited by a child process after `fork(2)`. The `memcntl()` function fails if it attempts to lock more memory than a system-specific limit.

Due to the potential impact on system resources, all operations, with the exception of MC_SYNC, are restricted to processes with super-user effective user ID.

USAGE The `memcntl()` function subsumes the operations of `plock(3C)` and `mctl(3B)`.

RETURN VALUES Upon successful completion, `memcntl()` returns 0; otherwise, it returns -1 and sets `errno` to indicate an error.

ERRORS The `memcntl()` function will fail if:

EAGAIN Some or all of the memory identified by the operation could not be locked when MC_LOCK or MC_LOCKAS was specified.

EBUSY Some or all of the addresses in the range [*addr*, *addr + len*) are locked and MC_SYNC with the MS_INVALIDATE option was specified.

EINVAL The *addr* argument specifies invalid selection criteria or is not a multiple of the page size as returned by `sysconf(3C)`; the *addr* and/or *len* argument does not have the value 0 when MC_LOCKAS or MC_UNLOCKAS is specified; or the *arg* argument is not valid for the function specified.

ENOMEM Some or all of the 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 and MC_LOCK, MC_LOCKAS, MC_UNLOCK, or MC_UNLOCKAS was specified.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO

`fork(2)`, `mmap(2)`, `mprotect(2)`, `mctl(3B)`, `mlock(3C)`, `mlockall(3C)`, `msync(3C)`, `plock(3C)`, `sysconf(3C)`, `attributes(5)`

NAME	mincore – determine residency of memory pages
SYNOPSIS	<pre>#include <sys/types.h> int mincore(caddr_t addr, size_t len, char *vec);</pre>
DESCRIPTION	<p>The mincore() function determines the residency of the memory pages in the address space covered by mappings in the range [<i>addr</i>, <i>addr</i> + <i>len</i>]. The status is returned as a character-per-page in the character array referenced by <i>vec</i> (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, and to 0 to indicate that it is not. The settings of other bits in each character are undefined and may contain other information in future implementations.</p> <p>Because the status of a page can change between the time mincore() checks and returns the information, returned information might be outdated. Only locked pages are guaranteed to remain in memory; see mlock(3C).</p>
RETURN VALUES	Upon successful completion, mincore() returns 0. Otherwise, -1 is returned and errno is set to indicate the error.
ERRORS	<p>The mincore() function will fail if:</p> <p>EFAULT The <i>vec</i> argument points to an illegal address.</p> <p>EINVAL The <i>addr</i> argument is not a multiple of the page size as returned by sysconf(3C), or the <i>len</i> argument has a value less than or equal to 0.</p> <p>ENOMEM Addresses in the range [<i>addr</i>, <i>addr</i> + <i>len</i>] are invalid for the address space of a process or specify one or more pages which are not mapped.</p>
SEE ALSO	mmap(2) , mlock(3C) , sysconf(3C)

NAME	mkdir – make a directory
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/stat.h> int mkdir(const char *path, mode_t mode);</pre>
DESCRIPTION	<p>The mkdir() function creates a new directory named by the path name pointed to by <i>path</i>. The mode of the new directory is initialized from <i>mode</i> (see chmod(2) for values of mode). The protection part of the <i>mode</i> argument is modified by the process's file creation mask (see umask(2)).</p> <p>The directory's owner ID is set to the process's effective user ID. The directory's group ID is set to the process's effective group ID, or if the <code>S_ISGID</code> bit is set in the parent directory, then the group ID of the directory is inherited from the parent. The <code>S_ISGID</code> bit of the new directory is inherited from the parent directory.</p> <p>If <i>path</i> is a symbolic link, it is not followed.</p> <p>The newly created directory is empty with the exception of entries for itself (.) and its parent directory (..).</p> <p>Upon successful completion, mkdir() marks for update the <code>st_atime</code>, <code>st_ctime</code> and <code>st_mtime</code> fields of the directory. Also, the <code>st_ctime</code> and <code>st_mtime</code> fields of the directory that contains the new entry are marked for update.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, no directory is created, and <code>errno</code> is set to indicate the error.
ERRORS	<p>The mkdir() function will fail if:</p> <p>EACCES Either a component of the path prefix denies search permission or write permission is denied on the parent directory of the directory to be created.</p> <p>EDQUOT The directory where the new file entry is being placed cannot be extended because the user's quota of disk blocks on that file system has been exhausted; the new directory cannot be created because the user's quota of disk blocks on that file system has been exhausted; or the user's quota of inodes on the file system where the file is being created has been exhausted.</p> <p>EEXIST The named file already exists.</p>

EFAULT	The <i>path</i> argument points to an illegal address.
EIO	An I/O error has occurred while accessing the file system.
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
EMLINK	The maximum number of links to the parent directory would be exceeded.
ENAMETOOLONG	The length of the <i>path</i> argument exceeds <code>PATH_MAX</code> , or the length of a <i>path</i> component exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.
ENOENT	A component of the path prefix does not exist or is a null pathname.
ENOLINK	The <i>path</i> argument points to a remote machine and the link to that machine is no longer active.
ENOSPC	No free space is available on the device containing the directory.
ENOTDIR	A component of the path prefix is not a directory.
EROFS	The path prefix resides on a read-only file system.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

`chmod(2)`, `mknod(2)`, `umask(2)`, `attributes(5)`, `stat(5)`

NAME	mknod – make a directory, or a special or ordinary file																																																									
SYNOPSIS	<pre>#include <sys/stat.h> int mknod(const char *path, mode_t mode, dev_t dev);</pre>																																																									
DESCRIPTION	<p>The mknod() function creates a new file named by the path name pointed to by <i>path</i>. The file type and permissions of the new file are initialized from <i>mode</i>.</p> <p>The file type is specified in <i>mode</i> by the <code>S_IFMT</code> bits, which must be set to one of the following values:</p> <table border="0"> <tr> <td><code>S_IFIFO</code></td> <td></td> <td>fifo special</td> </tr> <tr> <td><code>S_IFCHR</code></td> <td></td> <td>character special</td> </tr> <tr> <td><code>S_IFDIR</code></td> <td></td> <td>directory</td> </tr> <tr> <td><code>S_IFBLK</code></td> <td></td> <td>block special</td> </tr> <tr> <td><code>S_IFREG</code></td> <td></td> <td>ordinary file</td> </tr> </table> <p>The file access permissions are specified in <i>mode</i> by the <code>0007777</code> bits, and may be constructed by a bitwise OR operation of the following values:</p> <table border="0"> <tr> <td><code>S_ISUID</code></td> <td><code>04000</code></td> <td>Set user ID on execution.</td> </tr> <tr> <td><code>S_ISGID</code></td> <td><code>020#0</code></td> <td>Set group ID on execution if # is 7, 5, 3, or 1. Enable mandatory file/record locking if # is 6, 4, 2, or 0</td> </tr> <tr> <td><code>S_ISVTX</code></td> <td><code>01000</code></td> <td>Save text image after execution.</td> </tr> <tr> <td><code>S_IRWXU</code></td> <td><code>00700</code></td> <td>Read, write, execute by owner.</td> </tr> <tr> <td><code>S_IRUSR</code></td> <td><code>00400</code></td> <td>Read by owner.</td> </tr> <tr> <td><code>S_IWUSR</code></td> <td><code>00200</code></td> <td>Write by owner.</td> </tr> <tr> <td><code>S_IXUSR</code></td> <td><code>00100</code></td> <td>Execute (search if a directory) by owner.</td> </tr> <tr> <td><code>S_IRWXG</code></td> <td><code>00070</code></td> <td>Read, write, execute by group.</td> </tr> <tr> <td><code>S_IRGRP</code></td> <td><code>00040</code></td> <td>Read by group.</td> </tr> <tr> <td><code>S_IWGRP</code></td> <td><code>00020</code></td> <td>Write by group.</td> </tr> <tr> <td><code>S_IXGRP</code></td> <td><code>00010</code></td> <td>Execute by group.</td> </tr> <tr> <td><code>S_IRWXO</code></td> <td><code>00007</code></td> <td>Read, write, execute (search) by others.</td> </tr> <tr> <td><code>S_IROTH</code></td> <td><code>00004</code></td> <td>Read by others.</td> </tr> <tr> <td><code>S_IWOTH</code></td> <td><code>00002</code></td> <td>Write by others</td> </tr> </table>	<code>S_IFIFO</code>		fifo special	<code>S_IFCHR</code>		character special	<code>S_IFDIR</code>		directory	<code>S_IFBLK</code>		block special	<code>S_IFREG</code>		ordinary file	<code>S_ISUID</code>	<code>04000</code>	Set user ID on execution.	<code>S_ISGID</code>	<code>020#0</code>	Set group ID on execution if # is 7, 5, 3, or 1. Enable mandatory file/record locking if # is 6, 4, 2, or 0	<code>S_ISVTX</code>	<code>01000</code>	Save text image after execution.	<code>S_IRWXU</code>	<code>00700</code>	Read, write, execute by owner.	<code>S_IRUSR</code>	<code>00400</code>	Read by owner.	<code>S_IWUSR</code>	<code>00200</code>	Write by owner.	<code>S_IXUSR</code>	<code>00100</code>	Execute (search if a directory) by owner.	<code>S_IRWXG</code>	<code>00070</code>	Read, write, execute by group.	<code>S_IRGRP</code>	<code>00040</code>	Read by group.	<code>S_IWGRP</code>	<code>00020</code>	Write by group.	<code>S_IXGRP</code>	<code>00010</code>	Execute by group.	<code>S_IRWXO</code>	<code>00007</code>	Read, write, execute (search) by others.	<code>S_IROTH</code>	<code>00004</code>	Read by others.	<code>S_IWOTH</code>	<code>00002</code>	Write by others
<code>S_IFIFO</code>		fifo special																																																								
<code>S_IFCHR</code>		character special																																																								
<code>S_IFDIR</code>		directory																																																								
<code>S_IFBLK</code>		block special																																																								
<code>S_IFREG</code>		ordinary file																																																								
<code>S_ISUID</code>	<code>04000</code>	Set user ID on execution.																																																								
<code>S_ISGID</code>	<code>020#0</code>	Set group ID on execution if # is 7, 5, 3, or 1. Enable mandatory file/record locking if # is 6, 4, 2, or 0																																																								
<code>S_ISVTX</code>	<code>01000</code>	Save text image after execution.																																																								
<code>S_IRWXU</code>	<code>00700</code>	Read, write, execute by owner.																																																								
<code>S_IRUSR</code>	<code>00400</code>	Read by owner.																																																								
<code>S_IWUSR</code>	<code>00200</code>	Write by owner.																																																								
<code>S_IXUSR</code>	<code>00100</code>	Execute (search if a directory) by owner.																																																								
<code>S_IRWXG</code>	<code>00070</code>	Read, write, execute by group.																																																								
<code>S_IRGRP</code>	<code>00040</code>	Read by group.																																																								
<code>S_IWGRP</code>	<code>00020</code>	Write by group.																																																								
<code>S_IXGRP</code>	<code>00010</code>	Execute by group.																																																								
<code>S_IRWXO</code>	<code>00007</code>	Read, write, execute (search) by others.																																																								
<code>S_IROTH</code>	<code>00004</code>	Read by others.																																																								
<code>S_IWOTH</code>	<code>00002</code>	Write by others																																																								

S_IXOTH	00001	Execute by others.
S_ISVTX		On directories, restricted deletion flag.

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 the effective group ID of the process. However, if the `S_ISGID` bit is set in the parent directory, then the group ID of the file is inherited from the parent. If the group ID of the new file does not match the effective group ID or one of the supplementary group IDs, the `S_ISGID` bit is cleared.

The access permission bits of *mode* are modified by the process's file mode creation mask: all bits set in the process's file mode creation mask are cleared (see `umask(2)`). 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. See `makedev(3C)`.

If *path* is a symbolic link, it is not followed.

RETURN VALUES

Upon successful completion, `mknod()` returns 0. Otherwise, it returns `-1`, the new file is not created, and `errno` is set to indicate the error.

ERRORS

The `mknod()` function will fail if:

EACCES	A component of the path prefix denies search permission, or write permission is denied on the parent directory.
EDQUOT	The directory where the new file entry is being placed cannot be extended because the user's quota of disk blocks on that file system has been exhausted, or the user's quota of inodes on the file system where the file is being created has been exhausted.
EEXIST	The named file exists.
EFAULT	The <i>path</i> argument points to an illegal address.
EINTR	A signal was caught during the execution of the <code>mknod()</code> function.
EINVAL	An invalid argument exists.
EIO	An I/O error occurred while accessing the file system.

ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
ENAMETOOLONG	The length of the <i>path</i> argument exceeds <code>PATH_MAX</code> , or the length of a <i>path</i> component exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.
ENOENT	A component of the path prefix specified by <i>path</i> does not name an existing directory or <i>path</i> is an empty string.
ENOLINK	The <i>path</i> argument points to a remote machine and the link to that machine is no longer active.
ENOSPC	The directory that would contain the new file cannot be extended or the file system is out of file allocation resources.
ENOTDIR	A component of the path prefix is not a directory.
EPERM	The effective user of the calling process is not super-user.
EROFS	The directory in which the file is to be created is located on a read-only file system.
The mknod() function may fail if:	
ENAMETOOLONG	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>PATH_MAX</code> .

USAGE Normally, applications should use the `mkdir(2)` routine to make a directory, since the function `mknod()` may not establish directory entries for the directory itself (`.`) and the parent directory (`..`), and appropriate permissions are not required. Similarly, `mkfifo(3C)` should be used in place of `mknod()` in order to create FIFOs.

The `mknod()` function may be invoked only by a privileged user for file types other than FIFO special.

SEE ALSO `chmod(2)`, `creat(2)`, `exec(2)`, `mkdir(2)`, `open(2)`, `stat(2)`, `umask(2)`, `makedev(3C)`, `mkfifo(3C)`, `stat(5)`

NAME	mmap – map pages of memory				
SYNOPSIS	<pre>#include <sys/mman.h> void *mmap(void *addr, size_t len, int prot, int flags, int fildes, off_t off);</pre>				
DESCRIPTION	<p>The mmap() function establishes a mapping between a process's address space and a file or shared memory object. The format of the call is as follows:</p> <pre>pa = mmap(addr, len, prot, flags, fildes, off);</pre> <p>The mmap() function establishes a mapping between the address space of the process at an address <i>pa</i> for <i>len</i> bytes to the memory object represented by the file descriptor <i>fildes</i> at offset <i>off</i> for <i>len</i> bytes. The value of <i>pa</i> is a function of the <i>addr</i> argument and values of <i>flags</i>, further described below. A successful mmap() call returns <i>pa</i> as its result. The address range starting at <i>pa</i> and continuing for <i>len</i> bytes will be legitimate for the possible (not necessarily current) address space of the process. The range of bytes starting at <i>off</i> and continuing for <i>len</i> bytes will be legitimate for the possible (not necessarily current) offsets in the file or shared memory object represented by <i>fildes</i>.</p> <p>The mmap() function allows [<i>pa</i>, <i>pa + len</i>) to extend beyond the end of the object both at the time of the mmap() and while the mapping persists, such as when the file is created prior to the mmap() call and has no contents, or when the file is truncated. Any reference to addresses beyond the end of the object, however, will result in the delivery of a SIGBUS or SIGSEGV signal. The mmap() function cannot be used to implicitly extend the length of files.</p> <p>The mapping established by mmap() replaces any previous mappings for those whole pages containing any part of the address space of the process starting at <i>pa</i> and continuing for <i>len</i> bytes.</p> <p>If the size of the mapped file changes after the call to mmap() as a result of some other operation on the mapped file, the effect of references to portions of the mapped region that correspond to added or removed portions of the file is unspecified.</p> <p>The mmap() function is supported for regular files and shared memory objects. Support for any other type of file is unspecified.</p> <p>The <i>prot</i> argument determines whether read, write, execute, or some combination of accesses are permitted to the data being mapped. The <i>prot</i> argument should be either PROT_NONE or the bitwise inclusive OR of one or more of the other flags in the following table, defined in the header <code><sys/mman.h></code>.</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">PROT_READ</td> <td>Data can be read.</td> </tr> <tr> <td>PROT_WRITE</td> <td>Data can be written.</td> </tr> </table>	PROT_READ	Data can be read.	PROT_WRITE	Data can be written.
PROT_READ	Data can be read.				
PROT_WRITE	Data can be written.				

PROT_EXEC Data can be executed.

PROT_NONE Data cannot be accessed.

If an implementation of **mmap()** for a specific platform cannot support the combination of access types specified by *prot*, the call to **mmap()** fails. An implementation may permit accesses other than those specified by *prot*; however, the implementation will not permit a write to succeed where PROT_WRITE has not been set or permit any access where PROT_NONE alone has been set. Each platform-specific implementation of **mmap()** supports the following values of *prot*: PROT_NONE, PROT_READ, PROT_WRITE, and the inclusive OR of PROT_READ and PROT_WRITE. On some platforms, the PROT_WRITE protection option is implemented as PROT_READ|PROT_WRITE and PROT_EXEC as PROT_READ|PROT_EXEC. The file descriptor *fdes* is opened with read permission, regardless of the protection options specified. If PROT_WRITE is specified, the application must have opened the file descriptor *fdes* with write permission unless MAP_PRIVATE is specified in the *flags* argument as described below.

The *flags* argument provides other information about the handling of the mapped data. The value of *flags* is the bitwise inclusive OR of these options, defined in `<sys/mman.h>`:

MAP_SHARED Changes are shared.

MAP_PRIVATE Changes are private.

MAP_FIXED Interpret *addr* exactly.

MAP_NORESERVE Do not reserve swap space.

The MAP_SHARED and MAP_PRIVATE options describe the disposition of write references to the underlying 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 private copy is not created until the first write; until then, other users who have the object mapped MAP_SHARED can change the object. Either MAP_SHARED or MAP_PRIVATE must be specified, but not both. The mapping type is retained across **fork(2)**.

When MAP_FIXED is set in the *flags* argument, the system is informed that the value of *pa* must be *addr*, exactly. If MAP_FIXED is set, **mmap()** may return `(void *)-1` and set `errno` to `EINVAL`. If a MAP_FIXED request is successful, the mapping established by **mmap()** replaces any previous mappings for the process's pages in the range `[pa, pa + len)`. The use of MAP_FIXED is discouraged, since it may prevent a system from making the most effective use of its resources.

When `MAP_FIXED` is set and the requested address is the same as previous mapping, the previous address is unmapped and the new mapping is created on top of the old one.

When `MAP_FIXED` is not set, the system uses *addr* to arrive at *pa*. The *pa* so chosen will be an area of the address space that the system deems suitable for a mapping of *len* bytes to the file. The `mmap()` function interprets an *addr* value of 0 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 `MAP_NORESERVE` option specifies that no swap space be reserved for a mapping. Without this flag, the creation of a writable `MAP_PRIVATE` mapping reserves swap space equal to the size of the mapping; when the mapping is written into, the reserved space is employed to hold private copies of the data. A write into a `MAP_NORESERVE` mapping produces results which depend on the current availability of swap space in the system. If space is available, the write succeeds and a private copy of the written page is created; if space is not available, the write fails and a `SIGBUS` or `SIGSEGV` signal is delivered to the writing process. `MAP_NORESERVE` mappings are inherited across `fork()`; at the time of the `fork()`, swap space is reserved in the child for all private pages that currently exist in the parent; thereafter the child’s mapping behaves as described above.

The *off* argument is constrained to be aligned and sized according to the value returned by `sysconf(3C)` when passed `_SC_PAGESIZE` or `_SC_PAGE_SIZE`. When `MAP_FIXED` is specified, the *addr* argument must also meet these constraints. The system performs mapping operations over whole pages. Thus, while the *len* argument 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*).

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 the delivery of a `SIGBUS` or `SIGSEGV` signal. `SIGBUS` signals may also be delivered on various file system conditions, including quota exceeded errors.

The `mmap()` function adds an extra reference to the file associated with the file descriptor *fd* which is not removed by a subsequent `close(2)` on that file descriptor. This reference is removed when there are no more mappings to the file by a call to the `munmap(2)` function.

The `st_atime` field of the mapped file may be marked for update at any time between the `mmap()` call and the corresponding `munmap(2)` call. The initial read or write reference to a mapped region will cause the file's `st_atime` field to be marked for update if it has not already been marked for update.

The `st_ctime` and `st_mtime` fields of a file that is mapped with `MAP_SHARED` and `PROT_WRITE`, will be marked for update at some point in the interval between a write reference to the mapped region and the next call to `msync(3C)` with `MS_ASYNC` or `MS_SYNC` for that portion of the file by any process. If there is no such call, these fields may be marked for update at any time after a write reference if the underlying file is modified as a result.

If the process calls `mlockall(3C)` 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

Upon successful completion, the `mmap()` function returns the address at which the mapping was placed (*pa*); otherwise, it returns a value of `MAP_FAILED` and sets `errno` to indicate the error. The symbol `MAP_FAILED` is defined in the header `<sys/mman.h>`. No successful return from `mmap()` will return the value `MAP_FAILED`.

If `mmap()` fails for reasons other than `EBADF`, `EINVAL` or `ENOTSUP`, some of the mappings in the address range starting at *addr* and continuing for *len* bytes may have been unmapped.

ERRORS

The `mmap()` function will fail if:

- | | |
|---------------|--|
| EACCES | The <i>fil</i> des file descriptor is not open for read, regardless of the protection specified, or <i>fil</i> des is not open for write and <code>PROT_WRITE</code> was specified for a <code>MAP_SHARED</code> type mapping. |
| EAGAIN | The mapping could not be locked in memory.

There was insufficient room to reserve swap space for the mapping.

The file to be mapped is already locked using advisory or mandatory record locking. See <code>fcntl(2)</code> . |
| EBADF | The <i>fil</i> des file descriptor is not open. |
| EINVAL | The arguments <i>addr</i> (if <code>MAP_FIXED</code> was specified) or <i>off</i> are not multiples of the page size as returned by <code>sysconf()</code> .

The field in <i>flags</i> is invalid (neither <code>MAP_PRIVATE</code> or <code>MAP_SHARED</code> is set). |

	The argument <i>len</i> has a value less than or equal to 0.
EMFILE	The number of mapped regions would exceed an implementation-dependent limit (per process or per system).
ENODEV	The <i>fildev</i> argument refers to an object for which mmap() is meaningless, such as a terminal.
ENOMEM	<p>The <code>MAP_FIXED</code> option was specified and the range [<i>addr</i>, <i>addr + len</i>) exceeds that allowed for the address space of a process.</p> <p>The <code>MAP_FIXED</code> option was <i>not</i> specified and there is insufficient room in the address space to effect the mapping.</p> <p>The mapping could not be locked in memory, if required by <code>mlockall(3C)</code>, because it would require more space than the system is able to supply.</p> <p>The composite size of <i>len</i> plus the lengths obtained from all previous calls to mmap() exceeds <code>RLIMIT_VMEM</code> (see <code>getrlimit(2)</code>).</p>
ENOTSUP	The system does not support the combination of accesses requested in the <i>prot</i> argument.
ENXIO	<p>Addresses in the range [<i>off</i>, <i>off + len</i>) are invalid for the object specified by <i>fildev</i>.</p> <p>The <code>MAP_FIXED</code> option was specified in <i>flags</i> and the combination of <i>addr</i>, <i>len</i> and <i>off</i> is invalid for the object specified by <i>fildev</i>.</p>
E_OVERFLOW	The file is a regular file and the value of <i>off</i> plus <i>len</i> exceeds the offset maximum established in the open file description associated with <i>fildev</i> .
USAGE	<p>Use of mmap() may reduce the amount of memory available to other memory allocation functions.</p> <p>Use of <code>MAP_FIXED</code> may result in unspecified behaviour in further use of <code>brk(2)</code>, <code>sbrk(2)</code>, <code>malloc(3C)</code>, and <code>shmat(2)</code>. The use of <code>MAP_FIXED</code> is discouraged, as it may prevent an implementation from making the most effective use of resources.</p>

The application must ensure correct synchronization when using **mmap()** in conjunction with any other file access method, such as **read(2)** and **write(2)**, standard input/output, and **shmat(2)**.

The **mmap()** function has a transitional interface for 64-bit file offsets. See **1f64(5)**.

The **mmap()** function allows access to resources using address space manipulations instead of the **read()/write()** interface. Once a file is mapped, all a process has to do to access it is use the data at the address to which the object was mapped.

Consider the following pseudo-code:

```
fildev = open(...)  
lseek(fildev, offset, whence)  
read(fildev, buf, len)  
/* use data in buf */
```

The following is a rewrite using **mmap()**:

```
fildev = open(...)  
address = mmap((caddr_t) 0, len, (PROT_READ | PROT_WRITE),  
              MAP_PRIVATE, fildev, offset)  
/* use data at address */
```

SEE ALSO

close(2), **exec(2)**, **fcntl(2)**, **fork(2)**, **getrlimit(2)**, **mprotect(2)**, **munmap(2)**, **shmat(2)**, **lockf(3C)**, **mlockall(3C)**, **msync(3C)**, **plock(3C)**, **sysconf(3C)**, **1f64(5)**

NAME	mount – mount a file system
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/mount.h> int mount(const char *spec, const char *dir, int mflag, /* char *fstype, const char *dataptr, int datalen */ ...);</pre>
DESCRIPTION	<p>The mount() function requests that a removable file system contained on the block special file identified by <i>spec</i> be mounted on the directory identified by <i>dir</i>. The <i>spec</i> and <i>dir</i> arguments are pointers to path names. The <i>fstype</i> argument is the file system type, which can be determined by the sysfs(2) function. If both the <code>MS_DATA</code> and <code>MS_FSS</code> flag bits of <i>mflag</i> are off, the file system type defaults to the root file system type. Only if either flag is on is <i>fstype</i> used to indicate the file system type.</p> <p>If the <code>MS_DATA</code> flag is set in <i>mflag</i>, the system expects the <i>dataptr</i> and <i>datalen</i> arguments to be present. Together they describe a block of file system-specific data at address <i>dataptr</i> of length <i>datalen</i>. This is interpreted by file system-specific code within the operating system and its format depends on the file system type. If a particular file system type does not require this data, <i>dataptr</i> and <i>datalen</i> should both be zero. Note that <code>MS_FSS</code> is obsolete and is ignored if <code>MS_DATA</code> is also set, but if <code>MS_FSS</code> is set and <code>MS_DATA</code> is not, <i>dataptr</i> and <i>datalen</i> are both assumed to be 0.</p> <p>After a successful call to mount(), all references to the file <i>dir</i> refer to the root directory on the mounted file system.</p> <p>The low-order bit of <i>mflag</i> is used to control write permission on the mounted file system: if 1, writing is forbidden; otherwise writing is permitted according to individual file accessibility.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The mount() function will fail if:</p> <p>EBUSY The <i>dir</i> argument is currently mounted on, is someone's current working directory, or is otherwise busy; the device associated with <i>spec</i> is currently mounted; or there are no more mount table entries.</p> <p>EFAULT The <i>spec</i>, <i>dir</i>, or <i>datalen</i> argument points outside the allocated address space of the process.</p> <p>EINVAL The super block has an invalid magic number or the <i>fstype</i> is invalid.</p>

ELOOP	Too many symbolic links were encountered in translating <i>spec</i> or <i>dir</i> .
ENAMETOOLONG	The length of the <i>path</i> argument exceeds <code>PATH_MAX</code> , or the length of a <i>path</i> component exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.
ENOENT	None of the named files exists or is a null pathname.
ENOLINK	The <i>path</i> argument points to a remote machine and the link to that machine is no longer active.
ENOSPC	The file system state in the super-block is not <code>FSOKAY</code> and <i>mflag</i> requests write permission.
ENOTBLK	The <i>spec</i> argument is not a block special device.
ENOTDIR	The <i>dir</i> argument is not a directory, or a component of a path prefix is not a directory.
ENXIO	The device associated with <i>spec</i> does not exist.
EPERM	The effective user ID is not super-user.
EREMOTE	The <i>spec</i> argument is remote and cannot be mounted.
EROFS	The <i>spec</i> argument is write protected and <i>mflag</i> requests write permission.

USAGE The `mount()` function can be invoked only by processes with super-user privileges.

SEE ALSO `mount(1M)`, `sysfs(2)`, `umount(2)`

NAME	mprotect – set protection of memory mapping
SYNOPSIS	<pre>#include <sys/mman.h> int mprotect(void *addr, size_t len, int prot);</pre>
DESCRIPTION	<p>The mprotect() function changes the access protections on the mappings specified by the range [<i>addr</i>, <i>addr + len</i>), rounding <i>len</i> up to the next multiple of the page size as returned by sysconf(3C), to be that specified by <i>prot</i>. Legitimate values for <i>prot</i> are the same as those permitted for mmap(2) and are defined in <code><sys/mman.h></code> as:</p> <pre>PROT_READ /* page can be read */ PROT_WRITE /* page can be written */ PROT_EXEC /* page can be executed */ PROT_NONE /* page can not be accessed */</pre> <p>When mprotect() fails for reasons other than EINVAL, the protections on some of the pages in the range [<i>addr</i>, <i>addr + len</i>) may have been changed. If the error occurs on some page at <i>addr2</i>, then the protections of all whole pages in the range [<i>addr</i>, <i>addr2</i>] will have been modified.</p>
RETURN VALUES	Upon successful completion, mprotect() returns 0. Otherwise, it returns <code>-1</code> and sets <code>errno</code> to indicate the error.
ERRORS	<p>The mprotect() function will fail if:</p> <p>EACCES The <i>prot</i> argument specifies a protection that violates the access permission the process has to the underlying memory object.</p> <p>EINVAL The <i>len</i> argument has a value equal to 0, or <i>addr</i> is not a multiple of the page size as returned by sysconf(3C).</p> <p>ENOMEM 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.</p> <p>The mprotect() function may fail if:</p> <p>EAGAIN The address range [<i>addr</i>, <i>addr + len</i>) includes one or more pages that have been locked in memory and that were mapped <code>MAP_PRIVATE</code>; <i>prot</i> includes <code>PROT_WRITE</code>; and the system has insufficient resources to reserve memory for the private pages that may be created. These private pages may be created by store operations in the now-writable address range.</p>

SEE ALSO `mmap(2)`, `plock(3C)`, `mlock(3C)`, `mlockall(3C)`, `sysconf(3C)`

NAME	msgctl – message control operations
SYNOPSIS	<pre>#include <sys/msg.h> int msgctl(int msqid, int cmd, struct msgid_ds *buf);</pre>
DESCRIPTION	<p>The msgctl() function provides a variety of message control operations as specified by <i>cmd</i>. The following <i>cmds</i> are available:</p> <p>IPC_STAT Place the current value of each member of the data structure associated with <i>msqid</i> into the structure pointed to by <i>buf</i>. The contents of this structure are defined in intro(2).</p> <p>IPC_SET Set the value of the following members of the data structure associated with <i>msqid</i> to the corresponding value found in the structure pointed to by <i>buf</i>:</p> <pre>msg_perm.uid msg_perm.gid msg_perm.mode /* access permission bits only */ msg_qbytes</pre> <p>This <i>cmd</i> 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 <code>msg_perm.cuid</code> or <code>msg_perm.uid</code> in the data structure associated with <i>msqid</i>. Only super-user can raise the value of <code>msg_qbytes</code>.</p> <p>IPC_RMID Remove the message queue identifier specified by <i>msqid</i> from the system and destroy the message queue and data structure associated with it. This <i>cmd</i> 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 <code>msg_perm.cuid</code> or <code>msg_perm.uid</code> in the data structure associated with <i>msqid</i>. The <i>buf</i> argument is ignored.</p>
RETURN VALUES	Upon successful completion, msgctl() returns 0. Otherwise, it returns -1 and sets <code>errno</code> to indicate the error.
ERRORS	<p>The msgctl() function will fail if:</p> <p>EACCES The <i>cmd</i> argument is <code>IPC_STAT</code> and operation permission is denied to the calling process (see intro(2)).</p> <p>EFAULT The <i>buf</i> argument points to an illegal address.</p> <p>EINVAL The <i>msqid</i> argument is not a valid message queue identifier; or the <i>cmd</i> argument is not a valid command or is <code>IPC_SET</code> and <code>msg_perm.uid</code> or <code>msg_perm.gid</code> is not valid.</p>

- EPERM** The *cmd* argument is `IPC_RMID` or `IPC_SET` and the effective user ID of the calling process is not super-user and is not equal to the value of `msg_perm.cuid` or `msg_perm.uid` in the data structure associated with *msqid*.
- EPERM** The *cmd* argument is `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 super-user.
- E_OVERFLOW** The *cmd* argument is `IPC_STAT` and *uid* or *gid* is too large to be stored in the structure pointed to by *buf*.

SEE ALSO `intro(2)`, `msgget(2)`, `msgrcv(2)`, `msgsnd(2)`

NAME	msgget – get message queue
SYNOPSIS	<pre>#include <sys/msg.h> int msgget(key_t key, int msgflg);</pre>
DESCRIPTION	<p>The msgget() argument returns the message queue identifier associated with <i>key</i>.</p> <p>A message queue identifier and associated message queue and data structure (see intro(2)) are created for <i>key</i> if one of the following are true:</p> <ul style="list-style-type: none"> ■ <i>key</i> is <code>IPC_PRIVATE</code>. ■ <i>key</i> does not already have a message queue identifier associated with it, and $(msgflg \& IPC_CREAT)$ is true. <p>On creation, the data structure associated with the new message queue identifier is initialized as follows:</p> <ul style="list-style-type: none"> ■ <code>msg_perm.cuid</code>, <code>msg_perm.uid</code>, <code>msg_perm.cgid</code>, and <code>msg_perm.gid</code> are set to the effective user ID and effective group ID, respectively, of the calling process. ■ The low-order 9 bits of <code>msg_perm.mode</code> are set to the low-order 9 bits of <i>msgflg</i>. ■ <code>msg_qnum</code>, <code>msg_lspid</code>, <code>msg_lrpid</code>, <code>msg_stime</code>, and <code>msg_rtime</code> are set to 0. ■ <code>msg_ctime</code> is set to the current time. ■ <code>msg_qbytes</code> is set to the system limit.
RETURN VALUES	Upon successful completion, a non-negative integer representing a message queue identifier is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The msgget() function will fail if:</p> <p>EACCES A message queue identifier exists for <i>key</i>, but operation permission (see intro(2)) as specified by the low-order 9 bits of <i>msgflg</i> would not be granted.</p> <p>EEXIST A message queue identifier exists for <i>key</i> but $(msgflg \& IPC_CREAT)$ and $(msgflg \& IPC_EXCL)$ are both true.</p> <p>ENOENT A message queue identifier does not exist for <i>key</i> and $(msgflg \& IPC_CREAT)$ is false.</p>

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)`, `msgrcv(2)`, `msgsnd(2)`, `ftok(3C)`

NAME	msgrcv – message receive operation
SYNOPSIS	<pre>#include <sys/msg.h> int msgrcv(int msqid, void *msgp, size_t msgsz, long msgtyp, int msgflg);</pre>
DESCRIPTION	<p>The msgrcv() function reads a message from the queue associated with the message queue identifier specified by <i>msqid</i> and places it in the user-defined buffer pointed to by <i>msgp</i>.</p> <p>The <i>msgp</i> argument points to a user-defined buffer that must contain first a field of type <code>long int</code> that will specify the type of the message, and then a data portion that will hold the data bytes of the message. The structure below is an example of what this user-defined buffer might look like:</p> <pre>struct mymsg { long mtype; /* message type */ char mtext[1]; /* message text */ }</pre> <p>The <i>mtype</i> member is the received message's type as specified by the sending process.</p> <p>The <i>mtext</i> member is the text of the message.</p> <p>The <i>msgsz</i> argument specifies the size in bytes of <i>mtext</i>. The received message is truncated to <i>msgsz</i> bytes if it is larger than <i>msgsz</i> and (<i>msgflg</i>&MSG_NOERROR) is non-zero. The truncated part of the message is lost and no indication of the truncation is given to the calling process.</p> <p>The <i>msgtyp</i> argument specifies the type of message requested as follows:</p> <ul style="list-style-type: none"> ■ If <i>msgtyp</i> is 0, the first message on the queue is received. ■ If <i>msgtyp</i> is greater than 0, the first message of type <i>msgtyp</i> is received. ■ If <i>msgtyp</i> is less than 0, the first message of the lowest type that is less than or equal to the absolute value of <i>msgtyp</i> is received. <p>The <i>msgflg</i> argument specifies which of the following actions is to be taken if a message of the desired type is not on the queue:</p> <ul style="list-style-type: none"> ■ If (<i>msgflg</i>&IPC_NOWAIT) is non-zero, the calling process will return immediately with a return value of -1 and <i>errno</i> set to ENOMSG. ■ If (<i>msgflg</i>&IPC_NOWAIT) is 0, the calling process will suspend execution until one of the following occurs: <ul style="list-style-type: none"> ■ A message of the desired type is placed on the queue. ■ The message queue identifier <i>msqid</i> is removed from the system (see msgctl(2)); when this occurs, <i>errno</i> is set equal to EIDRM and -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 `sigaction(2)`.

Upon successful completion, the following actions are taken with respect to the data structure associated with `msgid` (see `intro(2)`):

- `msg_qnum` is decremented by 1.
- `msg_lrpid` is set equal to the process ID of the calling process.
- `msg_rtime` is set equal to the current time.

RETURN VALUES

Upon successful completion, `msgrcv()` returns a value equal to the number of bytes actually placed into the buffer `mtext`. Otherwise, `-1` is returned, no message is received, and `errno` is set to indicate the error.

ERRORS

The `msgrcv()` function will fail if:

- | | |
|---------------|---|
| E2BIG | The value of <code>mtext</code> is greater than <code>msgsz</code> and (<code>msgflg&MSG_NOERROR</code>) is 0. |
| EACCES | Operation permission is denied to the calling process. See <code>intro(2)</code> . |
| EIDRM | The message queue identifier <code>msgid</code> is removed from the system. |
| EINTR | The <code>msgrcv()</code> function was interrupted by a signal. |
| EINVAL | The <code>msgid</code> argument is not a valid message queue identifier; or the value of <code>msgsz</code> is less than 0. |
| ENOMSG | The queue does not contain a message of the desired type and (<code>msgflg&IPC_NOWAIT</code>) is non-zero. |

The `msgrcv()` function may fail if:

- | | |
|---------------|--|
| EFAULT | The <code>msgid</code> argument is not a valid message queue identifier. |
|---------------|--|

USAGE

The value passed as the `msgp` argument should be converted to type `void *`.

SEE ALSO

`intro(2)`, `msgctl(2)`, `msgget(2)`, `msgsnd(2)`, `sigaction(2)`

NAME	msgsnd – message send operation
SYNOPSIS	<pre>#include <sys/msg.h> int msgsnd(int msqid, const void *msgp, size_t msgsz, int msgflg);</pre>
DESCRIPTION	<p>The msgsnd() function is used to send a message to the queue associated with the message queue identifier specified by <i>msqid</i>.</p> <p>The <i>msgp</i> argument points to a user-defined buffer that must contain first a field of type <code>long int</code> that will specify the type of the message, and then a data portion that will hold the data bytes of the message. The structure below is an example of what this user-defined buffer might look like:</p> <pre>struct mymsg { long mtype; /* message type */ char mtext[1]; /* message text */ }</pre> <p>The <i>mtype</i> member is a non-zero positive type <code>long int</code> that can be used by the receiving process for message selection.</p> <p>The <i>mtext</i> member is any text of length <i>msgsz</i> bytes. The <i>msgsz</i> argument can range from 0 to a system-imposed maximum.</p> <p>The <i>msgflg</i> argument specifies the action to be taken if one or more of the following are true:</p> <ul style="list-style-type: none"> ■ The number of bytes already on the queue is equal to <code>msg_qbytes</code>; see intro(2). ■ The total number of messages on all queues system-wide is equal to the system-imposed limit. <p>These actions are as follows:</p> <ul style="list-style-type: none"> ■ If (<i>msgflg</i> & <code>IPC_NOWAIT</code>) is non-zero, the message will not be sent and the calling process will return immediately. ■ If (<i>msgflg</i> & <code>IPC_NOWAIT</code>) is 0, the calling process will suspend execution until one of the following occurs: <ul style="list-style-type: none"> ■ The condition responsible for the suspension no longer exists, in which case the message is sent. ■ The message queue identifier <i>msqid</i> is removed from the system (see msgctl(2)); when this occurs, <code>errno</code> is set equal to <code>EIDRM</code> and <code>-1</code> 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 `sigaction(2)`.

Upon successful completion, the following actions are taken with respect to the data structure associated with `msgid` (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.

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned, no message is sent, and `errno` is set to indicate the error.

ERRORS

The `msgsnd()` function will fail if:

- | | |
|---------------|---|
| EACCES | Operation permission is denied to the calling process. See <code>intro(2)</code> . |
| EAGAIN | The message cannot be sent for one of the reasons cited above and <code>(msgflg & IPC_NOWAIT)</code> is non-zero. |
| EIDRM | The message queue identifier <code>msgid</code> is removed from the system. |
| EINTR | The <code>msgsnd()</code> function was interrupted by a signal. |
| EINVAL | The value of <code>msgid</code> is not a valid message queue identifier, or the value of <code>mtype</code> is less than 1; or the value of <code>msgsz</code> is less than 0 or greater than the system-imposed limit. |

The `msgsnd()` function may fail if:

- | | |
|---------------|--|
| EFAULT | The <code>msgid</code> argument is not a valid message queue identifier. |
|---------------|--|

USAGE

The value passed as the `msgp` argument should be converted to type `void *`.

SEE ALSO

`intro(2)`, `msgctl(2)`, `msgget(2)`, `msgrcv(2)`, `sigaction(2)`

NAME	munmap – unmap pages of memory
SYNOPSIS	<pre>#include <sys/mman.h> int munmap(void *addr, size_t len);</pre>
DESCRIPTION	<p>The munmap() function removes the mappings for pages in the range <code>[addr, addr + len)</code>, rounding the <code>len</code> argument up to the next multiple of the page size as returned by sysconf(3C). If <code>addr</code> is not the address of a mapping established by a prior call to mmap(2), the behavior is undefined. After a successful call to munmap() and before any subsequent mapping of the unmapped pages, further references to these pages will result in the delivery of a SIGBUS or SIGSEGV signal to the process.</p> <p>The mmap(2) function often performs an implicit munmap().</p>
RETURN VALUES	Upon successful completion, munmap() returns 0; otherwise, it returns -1 and sets <code>errno</code> to indicate an error.
ERRORS	The munmap() function will fail if: EINVAL The <code>addr</code> argument is not a multiple of the page size as returned by sysconf(3C) ; addresses in the range <code>[addr, addr + len)</code> are outside the valid range for the address space of a process; or the <code>len</code> argument has a value less than or equal to 0.
SEE ALSO	mmap(2) , sysconf(3C)

NAME	nice – change priority of a process
SYNOPSIS	<pre>#include <unistd.h> int nice(int <i>incr</i>);</pre>
DESCRIPTION	<p>The nice() function allows a process to change its priority. The invoking process must be in a scheduling class that supports the nice().</p> <p>The nice() function adds the value of <i>incr</i> to the nice value of the calling process. A process's nice value is a non-negative number for which a greater positive value results in lower CPU priority.</p> <p>A maximum nice value of $(2 * NZERO) - 1$ and a minimum nice value of 0 are imposed by the system. NZERO is defined in <code><limits.h></code> with a default value of 20. Requests for values above or below these limits result in the nice value being set to the corresponding limit. A nice value of 40 is treated as 39.</p> <p>Only a process with super-user privileges can lower the nice value.</p>
RETURN VALUES	Upon successful completion, nice() returns the new nice value minus NZERO . Otherwise, <code>-1</code> is returned, the process's <i>nice</i> value is not changed, and <code>errno</code> is set to indicate the error.
ERRORS	The nice() function will fail if: EINVAL The nice() function is called by a process in a scheduling class other than time-sharing. EPERM The <i>inc</i> argument is negative or greater than 40 and the effective user ID of the calling process is not super-user.
USAGE	The <code>pricontrl(2)</code> function is a more general interface to scheduler functions. Since <code>-1</code> is a permissible return value in a successful situation, an application wishing to check for error situations should set <code>errno</code> to 0, then call nice() , and if it returns <code>-1</code> , check to see if <code>errno</code> is non-zero.
SEE ALSO	<code>nice(1)</code> , <code>exec(2)</code> , <code>pricontrl(2)</code>

NAME	ntp_adjtime – adjust local clock parameters
SYNOPSIS	<pre>#include <sys/timex.h> int ntp_adjtime(struct timex *tpr);</pre>
DESCRIPTION	<p>The ntp_adjtime() function adjusts the parameters used to discipline the local clock, according to the values in the struct <i>timex</i> pointed to by <i>tpr</i>. Before returning, it fills in the structure with the most recent values kept in the kernel.</p> <p>The adjustment is effected in part by speeding up or slowing down the clock, as necessary, and in part by phase-locking onto a once-per second pulse (PPS) provided by a driver, if available.</p> <pre>struct timex { uint32_t modes; /* clock mode bits (w) */ int32_t offset; /* time offset (us) (rw) */ int32_t freq; /* frequency offset (scaled ppm) (rw) */ int32_t maxerror; /* maximum error (us) (rw) */ int32_t esterror; /* estimated error (us) (rw) */ int32_t status; /* clock status bits (rw) */ int32_t constant; /* pll time constant (rw) */ int32_t precision; /* clock precision (us) (r) */ int32_t tolerance; /* clock frequency tolerance (scaled ppm) (r) */ int32_t ppsfreq; /* pps frequency (scaled ppm) (r) */ int32_t jitter; /* pps jitter (us) (r) */ int32_t shift; /* interval duration (s) (shift) (r) */ int32_t stabil; /* pps stability (scaled ppm) (r) */ int32_t jitcnt; /* jitter limit exceeded (r) */ int32_t calcnt; /* calibration intervals (r) */ int32_t errcnt; /* calibration errors (r) */ int32_t stbcnt; /* stability limit exceeded (r) */ };</pre>
RETURN VALUES	Upon successful completion, ntp_adjtime() returns the current clock state. Otherwise, it returns <code>IME_ERROR</code> .
ERRORS	<p>The ntp_adjtime() function will fail if:</p> <p>EFAULT The <i>tpr</i> argument is an invalid pointer.</p> <p>EPERM The user is not super-user.</p>
SEE ALSO	xntpd(1M) , ntp_gettime(2)

NAME	ntp_gettime – get local clock values
SYNOPSIS	<pre>#include <sys/timex.h> int ntp_gettime(struct ntptimeval *tpr);</pre>
DESCRIPTION	<p>The ntp_gettime() function reads the local clock value and dispersion, returning the information in <i>tpr</i>.</p> <p>The <code>ntptimeval</code> structure contains the following members:</p> <pre>struct ntptimeval { struct timeval time; /* current time (ro) */ int32_t maxerror; /* maximum error (us) (ro) */ int32_t esterror; /* estimated error (us) (ro) */ };</pre>
RETURN VALUES	Upon successful completion, ntp_gettime() returns 0. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The ntp_gettime() function will fail if:</p> <p>EFAULT The <i>tpr</i> argument points to an invalid address.</p> <p>The ntp_gettime() function will fail for 32-bit interfaces if:</p> <p>E_OVERFLOW The size of the <code>time.tv_sec</code> member of the <code>ntptimeval</code> structure pointed to by <i>tpr</i> is too small to contain the correct number of seconds.</p>
SEE ALSO	<code>xntpd(1M)</code> , <code>ntp_adjtime(2)</code>

NAME	open – open a file								
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/stat.h> #include <fcntl.h> int open(const char *path, int oflag, /* mode_t mode */...);</pre>								
DESCRIPTION	<p>The open() function establishes the connection between a file and a file descriptor. It creates an open file description that refers to a file and a file descriptor that refers to that open file description. The file descriptor is used by other I/O functions to refer to that file. The <i>path</i> argument points to a pathname naming the file.</p> <p>The open() function returns a file descriptor for the named file that is the lowest file descriptor not currently open for that process. The open file description is new, and therefore the file descriptor does not share it with any other process in the system. The <code>FD_CLOEXEC</code> file descriptor flag associated with the new file descriptor is cleared.</p> <p>The file offset used to mark the current position within the file is set to the beginning of the file.</p> <p>The file status flags and file access modes of the open file description is set according to the value of <i>oflag</i>.</p> <p>Values for <i>oflag</i> are constructed by a bitwise-inclusive-OR of flags from the following list, defined in <code><fcntl.h></code>. Applications must specify exactly one of the first three values (file access modes) below in the value of <i>oflag</i>:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>O_RDONLY</code></td> <td>Open for reading only.</td> </tr> <tr> <td style="padding-right: 20px;"><code>O_WRONLY</code></td> <td>Open for writing only.</td> </tr> <tr> <td style="padding-right: 20px;"><code>O_RDWR</code></td> <td>Open for reading and writing. The result is undefined if this flag is applied to a FIFO.</td> </tr> </table> <p>Any combination of the following may be used:</p> <table border="0"> <tr> <td style="padding-right: 40px;"><code>O_APPEND</code></td> <td>If set, the file offset is set to the end of the file prior to each write.</td> </tr> </table>	<code>O_RDONLY</code>	Open for reading only.	<code>O_WRONLY</code>	Open for writing only.	<code>O_RDWR</code>	Open for reading and writing. The result is undefined if this flag is applied to a FIFO.	<code>O_APPEND</code>	If set, the file offset is set to the end of the file prior to each write.
<code>O_RDONLY</code>	Open for reading only.								
<code>O_WRONLY</code>	Open for writing only.								
<code>O_RDWR</code>	Open for reading and writing. The result is undefined if this flag is applied to a FIFO.								
<code>O_APPEND</code>	If set, the file offset is set to the end of the file prior to each write.								

O_CREAT

If the file exists, this flag has no effect except as noted under O_EXCL below. Otherwise, the file is created with the user ID of the file set to the effective user ID of the process. The group ID of the file is set to the effective group IDs of the process, or if the S_ISGID bit is set in the directory in which the file is being created, the file's group ID is set to the group ID of its parent directory. If the group ID of the new file does not match the effective group ID or one of the supplementary groups IDs, the S_ISGID bit is cleared. The access permission bits (see <sys/stat.h>) of the file mode are set to the value of *mode*, modified as follows (see **creat(2)**): a bitwise-AND is performed on the file-mode bits and the corresponding bits in the complement of the process's file mode creation mask. Thus, all bits set in the process's file mode creation mask (see **umask(2)**) are correspondingly cleared in the file's permission mask. The "save text image after execution bit" of the mode is cleared (see **chmod(2)**).

O_SYNC Write I/O operations on the file descriptor complete as defined by synchronized I/O file integrity completion (see **fcntl(5)** definition of O_SYNC.) When bits other than the file permission bits are set, the effect is unspecified. The *mode* argument does not affect whether the file is open for reading, writing or for both.

O_DSYNC

Write I/O operations on the file descriptor complete as defined by synchronized I/O data integrity completion.

O_EXCL

If O_CREAT and O_EXCL are set, **open()** fails if the file exists. The check for the existence of the file and the creation of the file if it does not exist is atomic with respect to other processes executing **open()** naming the same filename in the same directory with O_EXCL and O_CREAT set. If O_CREAT is not set, the effect is undefined.

O_LARGEFILE

If set, the offset maximum in the open file description is the largest value that can be represented correctly in an object of type `off64_t`.

O_NOCTTY

If set and *path* identifies a terminal device, **open()** does not cause the terminal device to become the controlling terminal for the process.

O_NONBLOCK or O_NDELAY

These flags may affect subsequent reads and writes (see **read(2)** and **write(2)**). If both O_NDELAY and O_NONBLOCK are set, O_NONBLOCK takes precedence.

When opening a FIFO with O_RDONLY or O_WRONLY set:

If O_NONBLOCK or 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_NONBLOCK and O_NDELAY are clear:

An **open()** for reading only blocks until a process opens the file for writing. An **open()** for writing only blocks until a process opens the file for reading.

After both ends of a FIFO have been opened, there is no guarantee that further calls to **open()** O_RDONLY (O_WRONLY) will synchronize with later calls to **open()** O_WRONLY (O_RDONLY) until both ends of the

FIFO have been closed by all readers and writers. Any data written into a FIFO will be lost if both ends of the FIFO are closed before the data is read.

When opening a block special or character special file that supports non-blocking opens:

If `O_NONBLOCK` or `O_NDELAY` is set:

The **open()** function returns without blocking for the device to be ready or available. Subsequent behavior of the device is device-specific.

If `O_NONBLOCK` and `O_NDELAY` are clear:

The **open()** function blocks until the device is ready or available before returning.

Otherwise, the behavior of `O_NONBLOCK` and `O_NDELAY` is unspecified.

`O_RSYNC`

Read I/O operations on the file descriptor complete at the same level of integrity as specified by the `O_DSYNC` and `O_SYNC` flags. If both `O_DSYNC` and `O_RSYNC` are set in *oflag*, all I/O operations on the file descriptor complete as defined by synchronized I/O data integrity completion. If both `O_SYNC` and `O_RSYNC` are set in *oflag*, all I/O operations on the file descriptor complete as defined by synchronized I/O file integrity completion.

`O_SYNC`

If `O_SYNC` is set on a regular file, writes to that file cause the process to block until the data is delivered to the underlying hardware.

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 0 and the mode and owner are unchanged. It has no effect on FIFO special files or terminal device files. Its effect on other file types is implementation-dependent. The result of using O_TRUNC with O_RDONLY is undefined.

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 did previously exist, upon successful completion, **open()** marks for update the `st_ctime` and `st_mtime` fields of the file.

If *path* refers to a STREAMS file, *oflag* may be constructed from O_NONBLOCK or O_NODELAY OR-ed with either O_RDONLY, O_WRONLY, or O_RDWR. Other flag values are not applicable to STREAMS devices and have no effect on them. The values O_NONBLOCK and O_NODELAY affect the operation of STREAMS drivers and certain functions (see `read(2)`, `getmsg(2)`, `putmsg(2)`, and `write(2)`) applied to file descriptors associated with STREAMS files. For STREAMS drivers, the implementation of O_NONBLOCK and O_NODELAY is device-specific.

When **open()** is invoked to open a named stream, and the `connld` module (see `connld(7M)`) has been pushed on the pipe, **open()** blocks until the server process has issued an `I_RECVFD ioctl()` (see `streamio(7I)`) to receive the file descriptor.

If *path* names the master side of a pseudo-terminal device, then it is unspecified whether **open()** locks the slave side so that it cannot be opened. Portable applications must call `unlockpt(3C)` before opening the slave side.

If *path* is a symbolic link and O_CREAT and O_EXCL are set, the link is not followed.

Certain flag values can be set following **open()** as described in `fcntl(2)`.

The largest value that can be represented correctly in an object of type `off_t` is established as the offset maximum in the open file description.

RETURN VALUES

Upon successful completion, the **open()** function opens the file and return a non-negative integer representing the lowest numbered unused file descriptor. Otherwise, -1 is returned, `errno` is set to indicate the error, and no files are created or modified.

ERRORS	The open() function will fail if:
EACCES	Search permission is denied on a component of the path prefix, or the file exists and the permissions specified by <i>oflag</i> are denied, or the file does not exist and write permission is denied for the parent directory of the file to be created, or <code>O_TRUNC</code> is specified and write permission is denied.
EDQUOT	The file does not exist, <code>O_CREAT</code> is specified, and either the directory where the new file entry is being placed cannot be extended because the user's quota of disk blocks on that file system has been exhausted, or the user's quota of inodes on the file system where the file is being created has been exhausted.
EEXIST	The <code>O_CREAT</code> and <code>O_EXCL</code> flags are set, and the named file exists.
EINTR	A signal was caught during open() .
EFAULT	The <i>path</i> argument points to an illegal address.
EIO	The <i>path</i> argument names a STREAMS file and a hangup or error occurred during the open() .
EISDIR	The named file is a directory and <i>oflag</i> includes <code>O_WRONLY</code> or <code>O_RDWR</code> .
ELOOP	Too many symbolic links were encountered in resolving <i>path</i> .
EMFILE	<code>OPEN_MAX</code> file descriptors are currently open in the calling process.
EMULTIHOP	Components of <i>path</i> require hopping to multiple remote machines and the file system does not allow it.
ENAMETOOLONG	The length of the <i>path</i> argument exceeds <code>PATH_MAX</code> or a pathname component is longer than <code>NAME_MAX</code> .
ENFILE	The maximum allowable number of files is currently open in the system.
ENOENT	The <code>O_CREAT</code> flag is not set and the named file does not exist; or the <code>O_CREAT</code> flag is set and either the path prefix does not exist or the <i>path</i> argument points to an empty string.

ENOLINK	The <i>path</i> argument points to a remote machine, and the link to that machine is no longer active.
ENOSR	The <i>path</i> argument names a STREAMS-based file and the system is unable to allocate a STREAM.
ENOSPC	The directory or file system that would contain the new file cannot be expanded, the file does not exist, and <code>O_CREAT</code> is specified.
ENOTDIR	A component of the path prefix is not a directory.
ENXIO	The <code>O_NONBLOCK</code> flag is set, the named file is a FIFO, the <code>O_WRONLY</code> flag is set, and no process has the file open for reading; or the named file is a character special or block special file and the device associated with this special file does not exist.
EOPNOTSUPP	An attempt was made to open a path that corresponds to a <code>AF_UNIX</code> socket.
EOVERFLOW	The named file is a regular file and either <code>O_LARGEFILE</code> is not set and the size of the file cannot be represented correctly in an object of type <code>off_t</code> or <code>O_LARGEFILE</code> is set and the size of the file cannot be represented correctly in an object of type <code>off64_t</code> .
EROFS	The named file resides on a read-only file system and either <code>O_WRONLY</code> , <code>O_RDWR</code> , <code>O_CREAT</code> (if file does not exist), or <code>O_TRUNC</code> is set in the <i>oflag</i> argument.
The <code>open()</code> function may fail if:	
EAGAIN	The <i>path</i> argument names the slave side of a pseudo-terminal device that is locked.
EINVAL	The value of the <i>oflag</i> argument is not valid.
ENAMETOOLONG	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>PATH_MAX</code> .
ENOMEM	The <i>path</i> argument names a STREAMS file and the system is unable to allocate resources.
ETXTBSY	The file is a pure procedure (shared text) file that is being executed and <i>oflag</i> is <code>O_WRONLY</code> or <code>O_RDWR</code> .

USAGE The **open()** function has a transitional interface for 64-bit file offsets. See **lf64(5)**. Note that using **open64()** is equivalent to using **open()** with **O_LARGEFILE** set in *oflag*.

ATTRIBUTES See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO **intro(2)**, **chmod(2)**, **close(2)**, **creat(2)**, **dup(2)**, **exec(2)**, **fcntl(2)**, **getmsg(2)**, **getrlimit(2)**, **lseek(2)**, **putmsg(2)**, **read(2)**, **stat(2)**, **umask(2)**, **write(2)**, **unlockpt(3C)**, **attributes(5)**, **fcntl(5)**, **lf64(5)**, **stat(5)**, **conlnd(7M)**, **streamio(7I)**

NOTES Hierarchical Storage Management (HSM) file systems can sometimes cause long delays when opening a file, since HSM files must be recalled from secondary storage.

NAME	pause – suspend process until signal				
SYNOPSIS	#include <unistd.h> int pause(void);				
DESCRIPTION	The pause() function suspends the calling process until it receives a signal. The signal must be one that is not currently set to be ignored by the calling process. If the signal causes termination of the calling process, pause() does not return. If the signal is caught by the calling process and control is returned from the signal-catching function (see signal(3C)), the calling process resumes execution from the point of suspension.				
RETURN VALUES	Since pause() suspends thread execution indefinitely unless interrupted by a signal, there is no successful completion return value. If interrupted, it returns -1 and sets errno to indicate the error.				
ERRORS	The pause() function will fail if: EINTR A signal is caught by the calling process and control is returned from the signal-catching function.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	alarm(2) , kill(2) , wait(2) , signal(3C) , attributes(5)				

NAME pcsample – program execution time profile

SYNOPSIS #include <pcsample.h>

```
long pcsample(uintptr_t samples[], long nsamples);
```

DESCRIPTION The **pcsample()** function provides CPU-use statistics by profiling the amount of CPU time expended by a program.

For profiling dynamically-linked programs and 64-bit programs, it is superior to the **profil(2)** function, which assumes that the entire program is contained in a small, contiguous segment of the address space, divides this segment into “bins”, and on each clock tick increments the counter in the bin where the program is currently executing. With shared libraries creating discontinuous program segments spread throughout the address space, and with 64-bit address spaces so large that the size of “bins” would be measured in megabytes, the **profil()** function is of limited value.

The **pcsample()** function is passed an array *samples* containing *nsamples* pointer-sized elements. During program execution, the kernel samples the program counter of the process, storing unadulterated values in the array on each clock tick. The kernel stops writing to the array when it is full, which occurs after *nsamples* / HZ seconds of process virtual time. The HZ value is obtained by invoking the call `sysconf(_SC_CLK_TCK)`. See **sysconf(3C)**.

The sampling can be stopped by a subsequent call to **pcsample()** with the *nsamples* argument set to 0. Like **profil()**, sampling continues across a call to **fork(2)**, but is disabled by a call to one of the `exec` family of functions (see **exec(2)**). It is also disabled if an update of the *samples[]* array causes a memory fault.

RETURN VALUES The **pcsample()** function always returns 0 the first time it is called. On subsequent calls, it returns the number of samples that were stored during the previous invocation. If *nsamples* is invalid, it returns -1 and sets `errno` to indicate the error.

ERRORS The **pcsample()** function will fail if:

EINVAL The value of *nsamples* is not valid.

ATTRIBUTES See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe
Interface Stability	Stable

SEE ALSO `exec(2)`, `fork(2)`, `profil(2)`, `sysconf(3C)`, `attributes(5)`

NAME	pipe – create an interprocess channel				
SYNOPSIS	<pre>#include <unistd.h> int pipe(int fildes[2]);</pre>				
DESCRIPTION	<p>The pipe() function creates an I/O mechanism called a pipe and returns two file descriptors, <i>fildes</i>[0] and <i>fildes</i>[1]. The files associated with <i>fildes</i>[0] and <i>fildes</i>[1] are streams and are both opened for reading and writing. The <i>O_NDELAY</i> and <i>O_NONBLOCK</i> flags are cleared.</p> <p>A read from <i>fildes</i>[0] accesses the data written to <i>fildes</i>[1] on a first-in-first-out (FIFO) basis and a read from <i>fildes</i>[1] accesses the data written to <i>fildes</i>[0] also on a FIFO basis.</p> <p>The <i>FD_CLOEXEC</i> flag will be clear on both file descriptors.</p> <p>Upon successful completion pipe() marks for update the <i>st_atime</i>, <i>st_ctime</i>, and <i>st_mtime</i> fields of the pipe.</p>				
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <i>errno</i> is set to indicate the error.				
ERRORS	<p>The pipe() function will fail if:</p> <p>EMFILE There are <i>OPEN_MAX</i>-1 or more file descriptors currently open for this process.</p> <p>ENFILE A file table entry could not be allocated.</p>				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	sh(1) , fcntl(2) , fstat(2) , getmsg(2) , poll(2) , putmsg(2) , read(2) , write(2) , attributes(5) , streamio(7I)				
NOTES	Since a pipe is bi-directional, there are two separate flows of data. Therefore, the size (<i>st_size</i>) returned by a call to fstat(2) with argument <i>fildes</i> [0] or <i>fildes</i> [1] is the number of bytes available for reading from <i>fildes</i> [0] or <i>fildes</i> [1] respectively. Previously, the size (<i>st_size</i>) returned by a call to fstat() with argument <i>fildes</i> [1] (the write-end) was the number of bytes available for reading from <i>fildes</i> [0] (the read-end).				

NAME	poll – input/output multiplexing												
SYNOPSIS	<pre>#include <poll.h> int poll(struct pollfd fds[], nfds_t nfds, int timeout);</pre>												
DESCRIPTION	<p>The poll() function provides applications with a mechanism for multiplexing input/output over a set of file descriptors. For each member of the array pointed to by <i>fds</i>, poll() examines the given file descriptor for the event(s) specified in <i>events</i>. The number of <code>pollfd</code> structures in the <i>fds</i> array is specified by <i>nfds</i>. The poll() function identifies those file descriptors on which an application can read or write data, or on which certain events have occurred.</p> <p>The <i>fds</i> argument 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 member for each open file descriptor of interest. The array's members are <code>pollfd</code> structures, which contain the following members:</p> <pre>int fd; /* file descriptor */ short events; /* requested events */ short revents; /* returned events */</pre> <p>The <i>fd</i> member specifies an open file descriptor and the <i>events</i> and <i>revents</i> members are bitmasks constructed by a logical OR operation of any combination of the following event flags:</p> <table border="0"> <tr> <td style="padding-right: 20px;">POLLIN</td> <td>Data other than high priority data may be read without blocking. For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.</td> </tr> <tr> <td>POLLRDNORM</td> <td>Normal data (priority band equals 0) may be read without blocking. For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.</td> </tr> <tr> <td>POLLRDBAND</td> <td>Data from a non-zero priority band may be read without blocking. For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.</td> </tr> <tr> <td>POLLPRI</td> <td>High priority data may be received without blocking. For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.</td> </tr> <tr> <td>POLLOUT</td> <td>Normal data (priority band equals 0) may be written without blocking.</td> </tr> <tr> <td>POLLWRNORM</td> <td>The same as POLLOUT.</td> </tr> </table>	POLLIN	Data other than high priority data may be read without blocking. For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.	POLLRDNORM	Normal data (priority band equals 0) may be read without blocking. For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.	POLLRDBAND	Data from a non-zero priority band may be read without blocking. For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.	POLLPRI	High priority data may be received without blocking. For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.	POLLOUT	Normal data (priority band equals 0) may be written without blocking.	POLLWRNORM	The same as POLLOUT.
POLLIN	Data other than high priority data may be read without blocking. For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.												
POLLRDNORM	Normal data (priority band equals 0) may be read without blocking. For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.												
POLLRDBAND	Data from a non-zero priority band may be read without blocking. For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.												
POLLPRI	High priority data may be received without blocking. For STREAMS, this flag is set in <i>revents</i> even if the message is of zero length.												
POLLOUT	Normal data (priority band equals 0) may be written without blocking.												
POLLWRNORM	The same as POLLOUT.												

POLLWRBAND	Priority data (priority band > 0) may be written. This event only examines bands that have been written to at least once.
POLLERR	An error has occurred on the device or stream. This flag is only valid in the <code>revents</code> bitmask; it is not used in the <code>events</code> member.
POLLHUP	A hangup has occurred on the stream. This event and <code>POLLOUT</code> are mutually exclusive; a stream can never be writable if a hangup has occurred. However, this event and <code>POLLIN</code> , <code>POLLRDNORM</code> , <code>POLLRDBAND</code> , or <code>POLLPRI</code> are not mutually exclusive. This flag is only valid in the <code>revents</code> bitmask; it is not used in the <code>events</code> member.
POLLNVAL	The specified <code>fd</code> value does not belong to an open file. This flag is only valid in the <code>revents</code> member; it is not used in the <code>events</code> member.

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` member 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 are 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 `INFTIM` (or `-1`), `poll()` blocks until a requested event occurs or until the call is interrupted. The `poll()` function is not affected by the `O_NDELAY` and `O_NONBLOCK` flags.

The `poll()` function supports regular files, terminal and pseudo-terminal devices, STREAMS-based files, FIFOs and pipes. The behavior of `poll()` on elements of `fds` that refer to other types of file is unspecified.

The `poll()` function supports sockets.

A file descriptor for a socket that is listening for connections will indicate that it is ready for reading, once connections are available. A file descriptor for a socket that is connecting asynchronously will indicate that it is ready for writing, once a connection has been established.

	Regular files always poll <code>TRUE</code> for reading and writing.
RETURN VALUES	Upon successful completion, a non-negative value is returned. A positive value indicates the total number of file descriptors that has been selected (that is, file descriptors for which the <code>revents</code> member is non-zero). A value of 0 indicates that the call timed out and no file descriptors have been selected. Upon failure, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	The <code>poll()</code> function will fail if: EAGAIN Allocation of internal data structures failed, but the request may be attempted again. EFAULT Some argument points to an illegal address. EINTR A signal was caught during the <code>poll()</code> function. EINVAL The argument <code>nfds</code> is greater than <code>{OPEN_MAX}</code> , or one of the <code>fd</code> members refers to a <code>STREAM</code> or multiplexer that is linked (directly or indirectly) downstream from a multiplexer.
SEE ALSO	<code>intro(2)</code> , <code>getmsg(2)</code> , <code>getrlimit(2)</code> , <code>putmsg(2)</code> , <code>read(2)</code> , <code>write(2)</code> , <code>select(3C)</code> , <code>chpoll(9E)</code> <i>STREAMS Programming Guide</i>
NOTES	Non- <code>STREAMS</code> drivers use <code>chpoll(9E)</code> to implement <code>poll()</code> on these devices.

NAME	p_online – return or change processor operational status
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/processor.h> int p_online(processorid_t processorid, int flag);</pre>
DESCRIPTION	<p>The p_online() function changes or returns the operational status of processors. The state of the processor specified by the <i>processorid</i> argument is changed to the state represented by the <i>flag</i> argument.</p> <p>Legal values for <i>flag</i> are P_STATUS, P_ONLINE, P_OFFLINE, and P_NOINTR.</p> <p>When <i>flag</i> is P_STATUS, no processor status change occurs, but the current processor status is returned.</p> <p>The P_ONLINE, P_OFFLINE, and P_NOINTR values for <i>flag</i> refer to valid processor states. A processor in the P_ONLINE state is allowed to process LWPs (lightweight processes) and perform system activities. The processor is also interruptible by I/O devices attached to the system.</p> <p>A processor in the P_OFFLINE state is not allowed to process LWPs. The processor is as inactive as possible. If the hardware supports such a feature, the processor is not interruptible by attached I/O devices.</p> <p>A processor in the P_NOINTR state is allowed to process LWPs, but it is not interruptible by attached I/O devices. Typically, interrupts, when they occur are routed to other processors in the system. Not all systems support putting a processor into the P_NOINTR state. It is not permitted to put all the processors of a system into the P_NOINTR state. At least one processor must always be available to service system clock interrupts.</p> <p>Processor numbers are integers, greater than or equal to 0, and are defined by the hardware platform. Processor numbers are not necessarily contiguous, but “not too sparse.” Processor numbers should always be printed in decimal.</p> <p>The number of processors present can be determined by calling <code>sysconf(_SC_NPROCESSORS_CONF)</code>. The list of valid processor numbers can be determined by calling p_online() with <i>processorid</i> values starting at 0 until all processors have been found. The EINVAL error is returned for invalid processor numbers. See EXAMPLES below.</p>
RETURN VALUES	On successful completion, the value returned is the previous state of the processor, P_ONLINE, P_OFFLINE, P_NOINTR, or P_POWEROFF. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The p_online() function will fail if:</p> <p>EPERM The effective user of the calling process is not super-user.</p>

EINVAL	A non-existent processor ID was specified or <i>flag</i> was invalid.
EBUSY	The <i>flag</i> was <code>P_OFFLINE</code> and the specified processor is the only on-line processor, there are currently LWPs bound to the processor, or the processor performs some essential function that cannot be performed by another processor.
EBUSY	The <i>flag</i> was <code>P_NOINTR</code> and the specified processor is the only interruptible processor in the system, or it handles interrupts that cannot be handled by another processor.
EBUSY	The specified processor is powered off and cannot be powered on because some platform- specific resource is not available.
ENOTSUP	The specified processor is powered off, and the platform does not support power on of individual processors.

EXAMPLES

EXAMPLE 1 List the legal processor numbers.

The following code sample will list the legal processor numbers:

```
#include <sys/unistd.h>
#include <sys/processor.h>
#include <sys/types.h>
#include <stdio.h>
#include <errno.h>

int
main()
{
    processorid_t i;
    int status;
    int n = sysconf(_SC_NPROCESSORS_ONLN);
    for (i = 0; n > 0; i++) {
        status = p_online(i, P_STATUS);
        if (status == -1 && errno == EINVAL)
            continue;
        printf("processor %d present\n", i);
        n--;
    }
    return (0);
}
```

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO

psradm(1M), psrinfo(1M), processor_bind(2), processor_info(2),
pset_create(2), sysconf(3C), attributes(5)

NAME	priocntl – process scheduler control								
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/priocntl.h> #include <sys/rtpriocntl.h> #include <sys/tspriocntl.h> long priocntl(idtype_t idtype, id_t id, int cmd, /* arg */ ...);</pre>								
DESCRIPTION	<p>The priocntl() function provides for control over the scheduling of an active light weight process (LWP).</p> <p>LWPs fall into distinct classes with a separate scheduling policy applied to each class. The two classes currently supported are the realtime class and the time-sharing class. The characteristics of these classes are described under the corresponding headings below. The class attribute of an LWP is inherited across the fork(2) and _lwp_create(2) functions and the exec family of functions (see exec(2)). The priocntl() function can be used to dynamically change the class and other scheduling parameters associated with a running LWP or set of LWPs given the appropriate permissions as explained below.</p> <p>In the default configuration, a runnable realtime LWP runs before any other LWP. Therefore, inappropriate use of realtime LWP can have a dramatic negative impact on system performance.</p> <p>The priocntl() function provides an interface for specifying a process, set of processes or an LWP to which the function is to apply. The priocntlset(2) function provides the same functions as priocntl(), but allows a more general interface for specifying the set of LWPs to which the function is to apply.</p> <p>For priocntl(), the <i>idtype</i> and <i>id</i> arguments are used together to specify the set of LWPs. The interpretation of <i>id</i> depends on the value of <i>idtype</i>. The possible values for <i>idtype</i> and corresponding interpretations of <i>id</i> are as follows:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">P_LWPID</td> <td>The <i>id</i> argument is an LWP ID. The <i>priocntl</i> function applies to the LWP with the specified ID within the calling process.</td> </tr> <tr> <td>P_PID</td> <td>The <i>id</i> argument is a process ID specifying a single process. The priocntl() function applies to all LWPs currently associated with the specified process.</td> </tr> <tr> <td>P_PPID</td> <td>The <i>id</i> argument is a parent process ID. The priocntl() function applies to all LWPs currently associated with processes with the specified parent process ID.</td> </tr> <tr> <td>P_PGID</td> <td>The <i>id</i> argument is a process group ID. The priocntl() function applies to all LWPs currently associated with processes in the specified process group.</td> </tr> </table>	P_LWPID	The <i>id</i> argument is an LWP ID. The <i>priocntl</i> function applies to the LWP with the specified ID within the calling process.	P_PID	The <i>id</i> argument is a process ID specifying a single process. The priocntl() function applies to all LWPs currently associated with the specified process.	P_PPID	The <i>id</i> argument is a parent process ID. The priocntl() function applies to all LWPs currently associated with processes with the specified parent process ID.	P_PGID	The <i>id</i> argument is a process group ID. The priocntl() function applies to all LWPs currently associated with processes in the specified process group.
P_LWPID	The <i>id</i> argument is an LWP ID. The <i>priocntl</i> function applies to the LWP with the specified ID within the calling process.								
P_PID	The <i>id</i> argument is a process ID specifying a single process. The priocntl() function applies to all LWPs currently associated with the specified process.								
P_PPID	The <i>id</i> argument is a parent process ID. The priocntl() function applies to all LWPs currently associated with processes with the specified parent process ID.								
P_PGID	The <i>id</i> argument is a process group ID. The priocntl() function applies to all LWPs currently associated with processes in the specified process group.								

P_SID	The <i>id</i> argument is a session ID. The prioctl() function applies to all LWPs currently associated with processes in the specified session.
P_CID	The <i>id</i> argument is a class ID (returned by the prioctl() <code>PC_GETCID</code> command as explained below). The prioctl() function applies to all LWPs in the specified class.
P_UID	The <i>id</i> argument is a user ID. The prioctl() function applies to all LWPs with this effective user ID.
P_GID	The <i>id</i> argument is a group ID. The prioctl() function applies to all LWPs with this effective group ID.
P_ALL	The prioctl() function applies to all existing LWPs. The value of <i>id</i> is ignored. The permission restrictions described below still apply.

An *id* value of `P_MYID` can be used in conjunction with the *idtype* value to specify the calling LWP's LWP ID, parent process ID, process group ID, session ID, class ID, user ID, or group ID.

In order to change the scheduling parameters of an LWP (using the `PC_SETPARMS` command as explained below) the real or effective user ID of the LWP calling **prioctl()** must match the real or effective user ID of the receiving LWP or the effective user ID of the calling LWP must be super-user. These are the minimum permission requirements enforced for all classes. An individual class may impose additional permissions requirements when setting LWPs to that class and/or when setting class-specific scheduling parameters.

A special `sys` scheduling class exists for the purpose of scheduling the execution of certain special system processes (such as the swapper process). It is not possible to change the class of any LWP to `sys`. In addition, any processes in the `sys` class that are included in a specified set of processes are disregarded by **prioctl()**. For example, an *idtype* of `P_UID` and an *id* value of 0 would specify all processes with a user ID of 0 except processes in the `sys` class and (if changing the parameters using `PC_SETPARMS`) the `init(1M)` process.

The `init` process is a special case. In order for a **prioctl()** call to change the class or other scheduling parameters of the `init` process (process ID 1), it must be the only process specified by *idtype* and *id*. The `init` process may be assigned to any class configured on the system, but the time-sharing class is almost always the appropriate choice. (Other choices may be highly undesirable; see the *System Administration Guide, Volume I* for more information.)

The data type and value of *arg* are specific to the type of command specified by *cmd*.

A structure with the following members is used by the PC_GETCID and PC_GETCLINFO commands.

```
id_t   pc_cid;                /* Class id */
char   pc_clname[PC_CLNMSZ]; /* Class name */
int    pc_clinfo[PC_CLINFOSZ]; /* Class information */
```

The *pc_cid* member is a class ID returned by the **priocntl()** PC_GETCID command. The *pc_clname* member is a buffer of size PC_CLNMSZ (defined in `<sys/priocntl.h>`) used to hold the class name (RT for realtime or TS for time-sharing).

The *pc_clinfo* member is a buffer of size PC_CLINFOSZ (defined in `<sys/priocntl.h>`) used to return data describing the attributes of a specific class. The format of this data is class-specific and is described under the appropriate heading (REALTIME CLASS or TIME-SHARING CLASS) below.

A structure with the following elements is used by the PC_SETPARMS and PC_GETPARMS commands.

```
id_t   pc_cid;                /* LWP class */
int    pc_clparms[PC_CLPARMSZ]; /* Class-specific params */
```

The *pc_cid* member is a class ID (returned by the **priocntl()** PC_GETCID command). The special class ID PC_CLNULL can also be assigned to *pc_cid* when using the PC_GETPARMS command as explained below.

The *pc_clparms* buffer holds class-specific scheduling parameters. The format of this parameter data for a particular class is described under the appropriate heading below. PC_CLPARMSZ is the length of the *pc_clparms* buffer and is defined in `<sys/priocntl.h>`.

COMMANDS

Available **priocntl()** commands are:

PC_GETCID

Get class ID and class attributes for a specific class given class name. The *idtype* and *id* arguments are ignored. If *arg* is non-null, it points to a structure of type *pcinfo_t*. The *pc_clname* buffer contains the name of the class whose attributes you are getting.

On success, the class ID is returned in `pc_cid`, the class attributes are returned in the `pc_clinfo` buffer, and the **prioctl()** call returns the total number of classes configured in the system (including the `sys` class). If the class specified by `pc_clname` is invalid or is not currently configured the **prioctl()** call returns `-1` with `errno` set to `EINVAL`. The format of the attribute data returned for a given class is defined in the `<sys/rtprioctl.h>` or `<sys/tpsrioctl.h>` header and described under the appropriate heading below.

If *arg* is a null pointer, no attribute data is returned but the **prioctl()** call still returns the number of configured classes.

PC_GETCLINFO

Get class name and class attributes for a specific class given class ID. The *idtype* and *id* arguments are ignored. If *arg* is non-null, it points to a structure of type `pcinfo_t`. The `pc_cid` member is the class ID of the class whose attributes you are getting.

On success, the class name is returned in the `pc_clname` buffer, the class attributes are returned in the `pc_clinfo` buffer, and the **prioctl()** call returns the total number of classes configured in the system (including the `sys` class). The format of the attribute data returned for a given class is defined in the `<sys/rtprioctl.h>` or `<sys/tpsrioctl.h>` header file and described under the appropriate heading below.

If *arg* is a null pointer, no attribute data is returned but the **prioctl()** call still returns the number of configured classes.

PC_SETPARMS

Set the class and class-specific scheduling parameters of the specified LWP(s) associated with the specified process(es). When this command is used with the *idtype* of `P_LWPID`, it will set the class and class-specific scheduling parameters of the LWP. The *arg* argument points to a structure of type `pcparms_t`. The `pc_cid` member specifies the class you are setting and the `pc_clparms` buffer contains the class-specific parameters you are setting. The format of the class-specific parameter data is defined in the `<sys/rtprioctl.h>` or `<sys/tpsrioctl.h>` header and described under the appropriate class heading below.

When setting parameters for a set of LWPs, **prioctl()** acts on the LWPs in the set in an implementation-specific order. If **prioctl()** encounters an error for one or more of the target processes, it may or may not continue through the set of LWPs, depending on the nature of the error. If the error is related to permissions (`EPERM`), **prioctl()** continues through the LWP set, resetting

the parameters for all target LWPs for which the calling LWP has appropriate permissions. The **prioctl()** function then returns `-1` with `errno` set to `EPERM` to indicate that the operation failed for one or more of the target LWPs. If **prioctl()** encounters an error other than permissions, it does not continue through the set of target LWPs but returns the error immediately.

PC_GETPARMS

Get the class and/or class-specific scheduling parameters of an LWP. The *arg* member points to a structure of type `pcparms_t`.

If `pc_cid` specifies a configured class and a single LWP belonging to that class is specified by the *idtype* and *id* values or the `procset` structure, then the scheduling parameters of that LWP are returned in the `pc_clparms` buffer. If the LWP specified does not exist or does not belong to the specified class, the **prioctl()** call returns `-1` with `errno` set to `ESRCH`.

If `pc_cid` specifies a configured class and a set of LWPs is specified, the scheduling parameters of one of the specified LWP belonging to the specified class are returned in the `pc_clparms` buffer and the **prioctl()** call returns the process ID of the selected LWP. The criteria for selecting an LWP to return in this case is class dependent. If none of the specified LWPs exist or none of them belong to the specified class the **prioctl()** call returns `-1` with `errno` set to `ESRCH`.

If `pc_cid` is `PC_CLNULL` and a single LWP is specified the class of the specified LWP is returned in `pc_cid` and its scheduling parameters are returned in the `pc_clparms` buffer.

PC_ADMIN

This command provides functionality needed for the implementation of the **dispadm(1M)** utility. It is not intended for general use by other applications.

REALTIME CLASS

The realtime class provides a fixed priority preemptive scheduling policy for those LWPs requiring fast and deterministic response and absolute user/application control of scheduling priorities. If the realtime class is configured in the system it should have exclusive control of the highest range of scheduling priorities on the system. This ensures that a runnable realtime LWP is given CPU service before any LWP belonging to any other class.

The realtime class has a range of realtime priority (`rt_pri`) values that may be assigned to an LWP within the class. Real-time priorities range from 0 to *x*,

where the value of *x* is configurable and can be determined for a specific installation by using the **priocntl()** `PC_GETCID` or `PC_GETCLINFO` command.

The realtime scheduling policy is a fixed priority policy. The scheduling priority of a realtime LWP is never changed except as the result of an explicit request by the user/application to change the `rt_pri` value of the LWP.

For an LWP in the realtime class, the `rt_pri` value is, for all practical purposes, equivalent to the scheduling priority of the LWP. The `rt_pri` value completely determines the scheduling priority of a realtime LWP relative to other LWPs within its class. Numerically higher `rt_pri` values represent higher priorities. Since the realtime class controls the highest range of scheduling priorities in the system it is guaranteed that the runnable realtime LWP with the highest `rt_pri` value is always selected to run before any other LWPs in the system.

In addition to providing control over priority, **priocntl()** provides for control over the length of the time quantum allotted to the LWP in the realtime class. The time quantum value specifies the maximum amount of time an LWP may run assuming that it does not complete or enter a resource or event wait state (*sleep*). Note that if another LWP becomes runnable at a higher priority, the currently running LWP may be preempted before receiving its full time quantum.

The system's process scheduler keeps the runnable realtime LWPs on a set of scheduling queues. There is a separate queue for each configured realtime priority and all realtime LWPs with a given `rt_pri` value are kept together on the appropriate queue. The LWPs on a given queue are ordered in FIFO order (that is, the LWP at the front of the queue has been waiting longest for service and receives the CPU first). Real-time LWPs that wake up after sleeping, LWPs which change to the realtime class from some other class, LWPs which have used their full time quantum, and runnable LWPs whose priority is reset by **priocntl()** are all placed at the back of the appropriate queue for their priority. An LWP that is preempted by a higher priority LWP remains at the front of the queue (with whatever time is remaining in its time quantum) and runs before any other LWP at this priority. Following a **fork(2)** or **_lwp_create(2)** function call by a realtime LWP, the parent LWP continues to run while the child LWP (which inherits its parent's `rt_pri` value) is placed at the back of the queue.

A structure with the following members (defined in `<sys/rtpriocntl.h>`) defines the format used for the attribute data for the realtime class.

```
short    rt_maxpri;        /* Maximum realtime priority */
```

The **priocntl()** `PC_GETCID` and `PC_GETCLINFO` commands return realtime class attributes in the `pc_clinfo` buffer in this format.

The `rt_maxpri` member specifies the configured maximum `rt_pri` value for the realtime class (if `rt_maxpri` is `x`, the valid realtime priorities range from 0 to `x`).

A structure with the following members (defined in `<sys/rtpriocntl.h>`) defines the format used to specify the realtime class-specific scheduling parameters of an LWP.

```
short    rt_pri;          /* Real-Time priority */
uint_t   rt_tqsecs;     /* Seconds in time quantum */
int      rt_tqnsecs;    /* Additional nanoseconds in quantum */
```

When using the **priocntl()** `PC_SETPARMS` or `PC_GETPARMS` commands, if `pc_cid` specifies the realtime class, the data in the `pc_clparms` buffer is in this format.

The above commands can be used to set the realtime priority to the specified value or get the current `rt_pri` value. Setting the `rt_pri` value of an LWP that is currently running or runnable (not sleeping) causes the LWP to be placed at the back of the scheduling queue for the specified priority. The LWP is placed at the back of the appropriate queue regardless of whether the priority being set is different from the previous `rt_pri` value of the LWP. Note that a running LWP can voluntarily release the CPU and go to the back of the scheduling queue at the same priority by resetting its `rt_pri` value to its current realtime priority value. In order to change the time quantum of an LWP without setting the priority or affecting the LWP's position on the queue, the `rt_pri` member should be set to the special value `RT_NOCHANGE` (defined in `<sys/rtpriocntl.h>`). Specifying `RT_NOCHANGE` when changing the class of an LWP to realtime from some other class results in the realtime priority being set to 0.

For the **priocntl()** `PC_GETPARMS` command, if `pc_cid` specifies the realtime class and more than one realtime LWP is specified, the scheduling parameters of the realtime LWP with the highest `rt_pri` value among the specified LWPs are returned and the LWP ID of this LWP is returned by the **priocntl()** call. If there is more than one LWP sharing the highest priority, the one returned is implementation-dependent.

The `rt_tqsecs` and `rt_tqnsecs` members are used for getting or setting the time quantum associated with an LWP or group of LWPs. `rt_tqsecs` is the number of seconds in the time quantum and `rt_tqnsecs` is the number of

additional nanoseconds in the quantum. For example setting `rt_tqsecs` to 2 and `rt_tqnsecs` to 500,000,000 (decimal) would result in a time quantum of two and one-half seconds. Specifying a value of 1,000,000,000 or greater in the `rt_tqnsecs` member results in an error return with `errno` set to `EINVAL`. Although the resolution of the `tq_nsecs` member is very fine, the specified time quantum length is rounded up by the system to the next integral multiple of the system clock's resolution. The maximum time quantum that can be specified is implementation-specific and equal to `LONG_MAX` ticks (defined in `<limits.h>`). Requesting a quantum greater than this maximum results in an error return with `errno` set to `ERANGE` (although infinite quanta may be requested using a special value as explained below). Requesting a time quantum of 0 (setting both `rt_tqsecs` and `rt_tqnsecs` to 0) results in an error return with `errno` set to `EINVAL`.

The `rt_tqnsecs` member can also be set to one of the following special values (defined in `<sys/rtpriocntl.h>`), in which case the value of `rt_tqsecs` is ignored:

<code>RT_TQINF</code>	Set an infinite time quantum.
<code>RT_TQDEF</code>	Set the time quantum to the default for this priority (see <code>rt_dptbl(4)</code>).
<code>RT_NOCHANGE</code>	Do not set the time quantum. This value is useful when you wish to change the realtime priority of an LWP without affecting the time quantum. Specifying this value when changing the class of an LWP to realtime from some other class is equivalent to specifying <code>RT_TQDEF</code> .

In order to change the class of an LWP to realtime (from any other class) the LWP invoking `prctl()` must have super-user privileges. In order to change the priority or time quantum setting of a realtime LWP, the LWP invoking `prctl()` must have super-user privileges or must itself be a realtime LWP whose real or effective user ID matches the real or effective user ID of the target LWP.

The realtime priority and time quantum are inherited across `fork(2)` and the `exec` family of functions (see `exec(2)`).

TIME-SHARING CLASS

The time-sharing scheduling policy provides for a fair and effective allocation of the CPU resource among LWPs with varying CPU consumption characteristics. The objectives of the time-sharing policy are to provide good response time to interactive LWPs and good throughput to CPU-bound jobs while providing a degree of user/application control over scheduling.

The time-sharing class has a range of time-sharing user priority (see `ts_upri` below) values that may be assigned to LWPs within the class. A `ts_upri` value of 0 is defined as the default base priority for the time-sharing class.

User priorities range from $-x$ to $+x$ where the value of x is configurable and can be determined for a specific installation by using the **priocntl()** `PC_GETCID` or `PC_GETCLINFO` command.

The purpose of the user priority is to provide some degree of user/application control over the scheduling of LWPs in the time-sharing class. Raising or lowering the `ts_upri` value of an LWP in the time-sharing class raises or lowers the scheduling priority of the LWP. It is not guaranteed, however, that an LWP with a higher `ts_upri` value will run before one with a lower `ts_upri` value. This is because the `ts_upri` value is just one factor used to determine the scheduling priority of a time-sharing LWP. The system may dynamically adjust the internal scheduling priority of a time-sharing LWP based on other factors such as recent CPU usage.

In addition to the system-wide limits on user priority (returned by the `PC_GETCID` and `PC_GETCLINFO` commands) there is a per LWP user priority limit (see `ts_uprilim` below), which specifies the maximum `ts_upri` value that may be set for a given LWP; by default, `ts_uprilim` is 0.

A structure with the following members (defined in `<sys/tspriocntl.h>`) defines the format used for the attribute data for the time-sharing class.

```
short    ts_maxupri;    /* Limits of user priority range */
```

The **priocntl()** `PC_GETCID` and `PC_GETCLINFO` commands return time-sharing class attributes in the `pc_clinfo` buffer in this format.

`ts_maxupri` specifies the configured maximum user priority value for the time-sharing class. If `ts_maxupri` is x , the valid range for both user priorities and user priority limits is from $-x$ to $+x$.

A structure with the following members (defined in `<sys/tspriocntl.h>`) defines the format used to specify the time-sharing class-specific scheduling parameters of an LWP.

```
short    ts_uprilim;    /* Time-Sharing user priority limit */
short    ts_upri;      /* Time-Sharing user priority */
```

When using the **priocntl()** `PC_SETPARMS` or `PC_GETPARMS` commands, if `pc_cid` specifies the time-sharing class, the data in the `pc_clparms` buffer is in this format.

For the **priocntl()** `PC_GETPARMS` command, if `pc_cid` specifies the time-sharing class and more than one time-sharing LWP is specified, the scheduling parameters of the time-sharing LWP with the highest `ts_upri` value among the specified LWPs is returned and the LWP ID of this LWP is returned by the **priocntl()** call. If there is more than one LWP sharing the highest user priority, the one returned is implementation-dependent.

Any time-sharing LWP may lower its own `ts_uprilm` (or that of another LWP with the same user ID). Only a time-sharing LWP with super-user privileges may raise a `ts_uprilm`. When changing the class of an LWP to time-sharing from some other class, super-user privileges are required in order to set the initial `ts_uprilm` to a value greater than 0. Attempts by a non-super-user LWP to raise a `ts_uprilm` or set an initial `ts_uprilm` greater than 0 fail with a return value of -1 and `errno` set to `EPERM`.

Any time-sharing LWP may set its own `ts_upri` (or that of another LWP with the same user ID) to any value less than or equal to the LWP's `ts_uprilm`. Attempts to set the `ts_upri` above the `ts_uprilm` (and/or set the `ts_uprilm` below the `ts_upri`) result in the `ts_upri` being set equal to the `ts_uprilm`.

Either of the `ts_uprilm` or `ts_upri` members may be set to the special value `TS_NOCHANGE` (defined in `<sys/tspriocntl.h>`) in order to set one of the values without affecting the other. Specifying `TS_NOCHANGE` for the `ts_upri` when the `ts_uprilm` is being set to a value below the current `ts_upri` causes the `ts_upri` to be set equal to the `ts_uprilm` being set. Specifying `TS_NOCHANGE` for a parameter when changing the class of an LWP to time-sharing (from some other class) causes the parameter to be set to a default value. The default value for the `ts_uprilm` is 0 and the default for the `ts_upri` is to set it equal to the `ts_uprilm` which is being set.

The time-sharing user priority and user priority limit are inherited across `fork()` and the `exec` family of functions.

RETURN VALUES

Unless otherwise noted above, **priocntl()** returns a value of 0 on success. On failure, **priocntl()** returns -1 and sets `errno` to indicate the error.

ERRORS

The **priocntl()** function fails if:

- | | |
|---------------|---|
| EAGAIN | An attempt to change the class of an LWP failed because of insufficient resources other than memory (for example, class-specific kernel data structures). |
| EFAULT | One of the arguments points to an illegal address. |
| EINVAL | The argument <i>cmd</i> was invalid, an invalid or unconfigured class was specified, or one of the parameters specified was invalid. |

- ENOMEM** An attempt to change the class of an LWP failed because of insufficient memory.
- EPERM** The effective user of the calling LWP is not super-user.
- ERANGE** The requested time quantum is out of range.
- ESRCH** None of the specified LWPs exist.

SEE ALSO

`priocntl(1)`, `dispadm(1M)`, `init(1M)`, `_lwp_create(2)`, `exec(2)`, `fork(2)`, `nice(2)`, `priocntlset(2)`, `rt_dptbl(4)`

System Administration Guide, Volume I System Interface Guide

NAME	prioctlset – generalized process scheduler control
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/procset.h> #include <sys/prioctl.h> #include <sys/rtprioctl.h> #include <sys/tsprioctl.h> long prioctlset(procset_t *psp, int cmd, /* arg */ ...);</pre>
DESCRIPTION	<p>The prioctlset() function changes the scheduling properties of running processes. prioctlset() has the same functions as the prioctl() function, but a more general way of specifying the set of processes whose scheduling properties are to be changed.</p> <p><i>cmd</i> specifies the function to be performed. <i>arg</i> is a pointer to a structure whose type depends on <i>cmd</i>. See prioctl(2) for the valid values of <i>cmd</i> and the corresponding <i>arg</i> structures.</p> <p><i>psp</i> is a pointer to a procset structure, which prioctlset() uses to specify the set of processes whose scheduling properties are to be changed. The procset structure contains the following members:</p> <pre>idop_t p_op; /* operator connecting left/right sets */ idtype_t p_lidtype; /* left set ID type */ id_t p_lid; /* left set ID */ idtype_t p_ridtype; /* right set ID type */ id_t p_rid; /* right set ID */</pre> <p>The <i>p_lidtype</i> and <i>p_lid</i> members specify the ID type and ID of one (“left”) set of processes; the <i>p_ridtype</i> and <i>p_rid</i> members specify the ID type and ID of a second (“right”) set of processes. ID types and IDs are specified just as for the prioctl() function. The <i>p_op</i> member specifies the operation to be performed on the two sets of processes to get the set of processes the function is to apply to. The valid values for <i>p_op</i> and the processes they specify are:</p> <p>POP_DIFF Set difference: processes in left set and not in right set.</p> <p>POP_AND Set intersection: processes in both left and right sets.</p> <p>POP_OR Set union: processes in either left or right sets or both.</p> <p>POP_XOR Set exclusive-or: processes in left or right set but not in both.</p> <p>The following macro, which is defined in <code><procset.h></code>, offers a convenient way to initialize a procset structure:</p> <pre>#define setprocset(psp, op, ltype, lid, rtype, rid) \ (psp) => p_op = (op), \</pre>

```
(psp)⇒p_lidtype = (ltype), \
(psp)⇒p_lid     = (lid), \
(psp)⇒p_ridtype = (rtype), \
(psp)⇒p_rid     = (rid),
```

RETURN VALUES

Unless otherwise noted above, **priocntlset()** returns a value of 0 on success. **priocntlset()** returns -1 on failure and sets `errno` to indicate the error.

ERRORS

priocntlset() fails if one or more of the following are true :

- EAGAIN** An attempt to change the class of a process failed because of insufficient resources other than memory (for example, class-specific kernel data structures).
- EFAULT** One of the arguments points to an illegal address.
- EINVAL** The argument *cmd* was invalid, an invalid or unconfigured class was specified, or one of the parameters specified was invalid.
- ENOMEM** An attempt to change the class of a process failed because of insufficient memory.
- EPERM** The effective user of the calling process is not super-user.
- ERANGE** The requested time quantum is out of range.
- ESRCH** None of the specified processes exist.

SEE ALSO

`priocntl(1)`, `priocntl(2)`

NAME	processor_bind – bind LWPs to a processor
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/processor.h> #include <sys/procset.h> int processor_bind(idtype_t idtype, id_t id, processorid_t processorid, processorid_t *obind);</pre>
DESCRIPTION	<p>The processor_bind() function binds the LWP (lightweight process) or set of LWPs specified by <i>idtype</i> and <i>id</i> to the processor specified by <i>processorid</i>. If <i>obind</i> is not NULL, this function also sets the <code>processorid_t</code> variable pointed to by <i>obind</i> to the previous binding of one of the specified LWPs, or to <code>PBIND_NONE</code> if the selected LWP was not bound.</p> <p>If <i>idtype</i> is <code>P_PID</code>, the binding effects all LWPs of the process with process ID (PID) <i>id</i>.</p> <p>If <i>idtype</i> is <code>P_LWPID</code>, the binding effects the LWP of the current process with LWP ID <i>id</i>.</p> <p>If <i>id</i> is <code>P_MYID</code>, the specified LWP or process is the current one.</p> <p>If <i>processorid</i> is <code>PBIND_NONE</code>, the processor bindings of the specified LWPs are cleared.</p> <p>If <i>processorid</i> is <code>PBIND_QUERY</code>, the processor bindings are not changed.</p> <p>The effective user of the calling process must be superuser, or its real or effective user ID must match the real or effective user ID of the LWPs being bound. If the calling process does not have permission to change all of the specified LWPs, the bindings of the LWPs for which it does have permission will be changed even though an error is returned.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The processor_bind() function will fail if:</p> <p>ESRCH No processes or LWPs were found to match the criteria specified by <i>idtype</i> and <i>id</i>.</p> <p>EINVAL The specified processor is not on-line, or the <i>idtype</i> argument was not <code>P_PID</code> or <code>P_LWPID</code>.</p> <p>EFAULT The location pointed to by <i>obind</i> was not NULL and not writable by the user.</p>

EPERM The effective user of the calling process is not superuser, and its real or effective user ID does not match the real or effective user ID of one of the LWPs being bound.

SEE ALSO psradm(1M), psrinfo(1M), p_online(2), pset_bind(2), sysconf(3C)

NAME	processor_info – determine type and status of a processor
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/processor.h> int processor_info(processorid_t processorid, processor_info_t *infp);</pre>
DESCRIPTION	<p>The processor_info() function returns the status of the processor specified by <i>processorid</i> in the <code>processor_info_t</code> structure pointed to by <i>infp</i>.</p> <p>The structure <code>processor_info_t</code> contains the following members:</p> <pre>int pi_state; char pi_processor_type[PI_TYPELEN]; char pi_fputypes[PI_FPUTYPE]; int pi_clock;</pre> <p>The <code>pi_state</code> member is the current state of the processor, either <code>P_ONLINE</code>, <code>P_OFFLINE</code>, or <code>P_POWEROFF</code>.</p> <p>The <code>pi_processor_type</code> member is a null-terminated ASCII string specifying the type of the processor.</p> <p>The <code>pi_fputypes</code> member is a null-terminated ASCII string containing the comma-separated types of floating-point units (FPUs) attached to the processor. This string will be empty if no FPU is attached.</p> <p>The <code>pi_clock</code> member is the processor clock frequency rounded to the nearest megahertz. It may be 0 if not known.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The processor_info() function will fail if:</p> <p>EINVAL An non-existent processor ID was specified.</p> <p>EFAULT The <code>processor_info_t</code> structure pointed to by <i>infp</i> was not writable by the user.</p>
SEE ALSO	<code>psradm(1M)</code> , <code>psrinfo(1M)</code> , <code>p_online(2)</code> , <code>sysconf(3C)</code>

NAME	profil – execution time profile
SYNOPSIS	<pre>#include <unistd.h> void profil(unsigned short *<i>buff</i>, unsigned int <i>bufsiz</i>, unsigned int <i>offset</i>, unsigned int <i>scale</i>);</pre>
DESCRIPTION	<p>The profil() function provides CPU-use statistics by profiling the amount of CPU time expended by a program. The profil() function generates the statistics by creating an execution histogram for a current process. The histogram is defined for a specific region of program code to be profiled, and the identified region is logically broken up into a set of equal size subdivisions, each of which corresponds to a count in the histogram. With each clock tick, the current subdivision is identified and its corresponding histogram count is incremented. These counts establish a relative measure of how much time is being spent in each code subdivision. The resulting histogram counts for a profiled region can be used to identify those functions that consume a disproportionately high percentage of CPU time.</p> <p>The <i>buff</i> argument is a buffer of <i>bufsiz</i> bytes in which the histogram counts are stored in an array of unsigned short int.</p> <p>The <i>offset</i>, <i>scale</i>, and <i>bufsiz</i> arguments specify the region to be profiled.</p> <p>The <i>offset</i> argument is effectively the start address of the region to be profiled.</p> <p>The <i>scale</i> argument is a contraction factor that indicates how much smaller the histogram buffer is than the region to be profiled. More precisely, <i>scale</i> is interpreted as an unsigned 16-bit fixed-point fraction with the decimal point implied on the left. Its value is the reciprocal of the number of bytes in a subdivision, per byte of histogram buffer. Since there are two bytes per histogram counter, the effective ratio of subdivision bytes per counter is one half the <i>scale</i>.</p> <p>The values of <i>scale</i> are as follows:</p> <ul style="list-style-type: none"> ■ the maximum value of <i>scale</i>, 0xffff (approximately 1), maps subdivisions 2 bytes long to each counter. ■ the minimum value of <i>scale</i> (for which profiling is performed), 0x0002 (1/32,768), maps subdivision 65,536 bytes long to each counter. ■ the default value of <i>scale</i> (currently used by <code>cc -cp</code>), 0x4000, maps subdivisions 8 bytes long to each counter. <p>The values are used within the kernel as follows: when the process is interrupted for a clock tick, the value of <i>offset</i> is subtracted from the current value of the program counter (<i>pc</i>), and the remainder is multiplied by <i>scale</i> to derive a result. That result is used as an index into the histogram array to</p>

locate the cell to be incremented. Therefore, the cell count represents the number of times that the process was executing code in the subdivision associated with that cell when the process was interrupted.

The value of *scale* can be computed as $(RATIO * 02000001)$, where *RATIO* is the desired ratio of *bufsiz* to profiled region size, and has a value between 0 and 1. Qualitatively speaking, the closer *RATIO* is to 1, the higher the resolution of the profile information.

The value of *bufsiz* can be computed as $(size_of_region_to_be_profiled * RATIO)$.

Profiling is turned off by giving a *scale* value of 0 or 1, and is rendered ineffective by giving a *bufsiz* value of 0. Profiling is turned off when one of the *exec* family of functions (see *exec(2)*) is executed, but remains on in both child and parent processes after a *fork(2)*. Profiling is turned off if a *buff* update would cause a memory fault.

USAGE

The *pcsample(2)* function should be used when profiling dynamically-linked programs and 64-bit programs.

SEE ALSO

exec(2), *fork(2)*, *pcsample(2)*, *times(2)*, *monitor(3C)*, *prof(5)*

NOTES

In Solaris releases prior to 2.6, calling *profil()* in a multithreaded program would impact only the calling LWP; the profile state was not inherited at LWP creation time. To profile a multithreaded program with a global profile buffer, each thread needed to issue a call to *profil()* at threads start-up time, and each thread had to be a bound thread. This was cumbersome and did not easily support dynamically turning profiling on and off. In Solaris 2.6, the *profil()* system call for multithreaded processes has global impact — that is, a call to *profil()* impacts all LWPs/threads in the process. This may cause applications that depend on the previous per-LWP semantic to break, but it is expected to improve multithreaded programs that wish to turn profiling on and off dynamically at runtime.

NAME	pset_bind – bind LWPs to a set of processors
SYNOPSIS	<pre>#include <sys/pset.h> int pset_bind(psetid_t pset, idtype_t idtype, id_t id, psetid_t *opset);</pre>
DESCRIPTION	<p>The pset_bind() function binds the LWP or set of LWPs specified by <i>idtype</i> and <i>id</i> to the processor set specified by <i>pset</i>. If <i>obind</i> is not <code>NULL</code>, pset_bind() sets the <code>psetid_t</code> variable pointed to by <i>opset</i> to the previous processor set binding of one of the specified LWP, or to <code>PS_NONE</code> if the selected LWP was not bound.</p> <p>If <i>idtype</i> is <code>P_PID</code>, the binding affects all LWPs of the process with process ID (PID) <i>id</i>.</p> <p>If <i>idtype</i> is <code>P_LWPID</code>, the binding affects the LWP of the current process with LWP ID <i>id</i>.</p> <p>If <i>id</i> is <code>P_MYID</code>, the specified LWP or process is the current one.</p> <p>If <i>pset</i> is <code>PS_NONE</code>, the processor set bindings of the specified LWPs are cleared.</p> <p>If <i>pset</i> is <code>PS_QUERY</code>, the processor set bindings are not changed.</p> <p>The effective user of the calling process must be super-user, or its real or effective user ID must match the real or effective user ID of the LWPs being bound, or <i>pset</i> must be <code>PS_QUERY</code>. If the calling process does not have permission to change all of the specified LWPs, the bindings of the LWPs for which it does have permission will be changed even though an error is returned.</p> <p>If the processor set type of <i>pset</i> is <code>PS_PRIVATE</code> (see pset_info(2)), the effective user of the calling process must be super-user.</p> <p>LWPs that have been bound to a processor with processor_bind(2) may also be bound to a processor set if the processor is part of the processor set. If this occurs, the binding to the processor remains in effect. If the processor binding is later removed, the processor set binding becomes effective.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The pset_bind() function will fail if:</p> <p>EBUSY One of the LWPs is bound to a processor, and the specified processor set does not include that processor.</p> <p>EFAULT The location pointed to by <i>opset</i> was not <code>NULL</code> and not writable by the user.</p>

- EINVAL** An invalid processor set ID was specified; or *idtype* was not P_PID or P_LWPID.
- EPERM** The effective user of the calling process is not super-user, and either the processor set type of *pset* is PS_USER, or the real or effective user ID of the calling process does not match the real or effective user ID of one of the LWPs being bound.
- ESRCH** No processes or LWPs were found to match the criteria specified by *idtype* and *id*.

SEE ALSO

`pbind(1M)`, `psrset(1M)`, `processor_bind(2)`, `pset_create(2)`,
`pset_info(2)`

NAME	pset_create, pset_destroy, pset_assign – manage sets of processors
SYNOPSIS	<pre>#include <sys/pset.h> int pset_create(psetid_t * newpset); int pset_destroy(psetid_t pset); int pset_assign(psetid_t pset, processorid_t cpu, psetid_t * opset);</pre>
DESCRIPTION	<p>These functions control the creation and management of sets of processors. Processor sets allow a subset of the system's processors to be set aside for exclusive use by specified LWP s and processes. The binding of LWP s and processes to processor sets is controlled by pset_bind(2) .</p> <p>The pset_create() function creates an empty processor set that contains no processors. On successful return, <i>newpset</i> will contain the ID of the new processor set.</p> <p>The pset_destroy() function destroys the processor set <i>pset</i> , releasing its constituent processors and processes.</p> <p>The pset_assign() function assigns the processor <i>cpu</i> to the processor set <i>pset</i> . A processor that has been assigned to a processor set will run only LWP s and processes that have been explicitly bound to that processor set, unless another LWP requires a resource that is only available on that processor. On successful return, if <i>opset</i> is non-null, <i>opset</i> will contain the processor set ID of the former processor set of the processor.</p> <p>If <i>pset</i> is <code>PS_NONE</code> , pset_assign() releases processor <i>cpu</i> from its current processor set.</p> <p>If <i>pset</i> is <code>PS_QUERY</code> , pset_assign() makes no change to processor sets, but returns the current processor set ID of processor <i>cpu</i> in <i>opset</i> .</p> <p>These functions are restricted to super-user use, except for pset_assign() when <i>pset</i> is <code>PS_QUERY</code> .</p>
RETURN VALUES	Upon successful completion, these functions return 0 . Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>These functions will fail if:</p> <p>EBUSY The processor could not be moved to the specified processor set.</p> <p>EFAULT The location pointed to by <i>newpset</i> was not writable by the user, or the location pointed to by <i>opset</i> was not <code>NULL</code> and not writable by the user.</p>

- EINVAL** The specified processor does not exist, the specified processor is not on-line, or an invalid processor set was specified.
- ENOMEM** There was insufficient space for `pset_create` to create a new processor set.
- EPERM** The effective user of the calling process is not super-user.

SEE ALSO

`psradm(1M)` , `psrinfo(1M)` , `psrset(1M)` , `p_online(2)` ,
`processor_bind(2)` , `pset_bind (2)` , `pset_info(2)`

NOTES

Processors belonging to different processor sets of type `PS_SYSTEM` (see `pset_info(2)`) cannot be assigned to the same processor set of type `PS_PRIVATE` . If this is attempted, `pset_assign()` will fail and set `errno` to `EINVAL` .

Processors with LWP s bound to them using `processor_bind(2)` cannot be assigned to a new processor set. If this is attempted, `pset_assign()` will fail and set `errno` to `EBUSY` .

NAME	pset_info – get information about a processor set
SYNOPSIS	<pre>#include <sys/pset.h> int pset_info(psetid_t pset, int *type, uint_t *numcpus, processorid_t *cpulist);</pre>
DESCRIPTION	<p>The pset_info() function returns information on the processor set <i>pset</i>.</p> <p>If <i>type</i> is non-null, then on successful completion the type of the processor set will be stored in the location pointed to by <i>type</i>. Processor set types can have the following values:</p> <p>PS_SYSTEM The processor set was created by the system. Processor sets of this type cannot be modified or removed by the user, but LWPs and processes can be bound to them using pset_bind(2).</p> <p>PS_PRIVATE The processor set was created by pset_create(2) and can be modified by pset_assign(2) and removed by pset_destroy(2). LWPs and processes can also be bound to this processor set using pset_bind().</p> <p>If <i>numcpus</i> is non-null, then on successful completion the number of processors in the processor set will be stored in the location pointed to by <i>numcpus</i>.</p> <p>If <i>numcpus</i> and <i>cpulist</i> are both non-null, then <i>cpulist</i> points to a buffer where a list of processors assigned to the processor set is to be stored, and <i>numcpus</i> points to the maximum number of processor IDs the buffer can hold. On successful completion, the list of processors up to the maximum buffer size is stored in the buffer pointed to by <i>cpulist</i>.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The pset_info() function will fail if:</p> <p>EFAULT The location pointed to by <i>type</i>, <i>numcpus</i>, or <i>cpulist</i> was not null and not writable by the user.</p> <p>EINVAL An invalid processor set ID was specified.</p>
SEE ALSO	psrinfo(1M) , psrset(1M) , processor_info(2) , pset_assign(2) , pset_bind(2) , pset_create(2) , pset_destroy(2)

NAME	ptrace – allows a parent process to control the execution of a child process
SYNOPSIS	<pre>#include <unistd.h> #include <sys/types.h> int ptrace(int request, pid_t pid, int addr, int data);</pre>
DESCRIPTION	<p>The ptrace() function allows a parent process to control the execution of a child process. Its primary use is for the implementation of breakpoint debugging. The child process behaves normally until it encounters a signal (see signal(5)), at which time it enters a stopped state and its parent is notified via the wait(2) function. When the child is in the stopped state, its parent can examine and modify its “core image” using ptrace(). Also, the parent can cause the child either to terminate or continue, with the possibility of ignoring the signal that caused it to stop.</p> <p>The <i>request</i> argument determines the action to be taken by ptrace() and is one of the following:</p> <ul style="list-style-type: none"> 0 This request must be issued by the child process if it is to be traced by its parent. It turns on the child’s trace flag that stipulates that the child should be left in a stopped state on receipt of a signal rather than the state specified by <i>func</i> (see signal(3C)). The <i>pid</i>, <i>addr</i>, and <i>data</i> arguments are ignored, and a return value is not defined for this request. Peculiar results ensue if the parent does not expect to trace the child. <p>The remainder of the requests can only be used by the parent process. For each, <i>pid</i> is the process ID of the child. The child must be in a stopped state before these requests are made.</p> <ul style="list-style-type: none"> 1, 2 With these requests, the word at location <i>addr</i> in the address space of the child is returned to the parent process. If instruction and data space are separated, request 1 returns a word from instruction space, and request 2 returns a word from data space. If instruction and data space are not separated, either request 1 or request 2 may be used with equal results. The <i>data</i> argument is ignored. These two requests fail if <i>addr</i> is not the start address of a word, in which case -1 is returned to the parent process and the parent’s <code>errno</code> is set to <code>EIO</code>. 3 With this request, the word at location <i>addr</i> in the child’s user area in the system’s address space (see <code><sys/user.h></code>) is returned to the parent process. The <i>data</i> argument is ignored. This request fails if <i>addr</i> is not the start address of a word or is outside the user area, in which case -1 is returned to the parent process and the parent’s <code>errno</code> is set to <code>EIO</code>. 4, 5 With these requests, the value given by the <i>data</i> argument is written into the address space of the child at location <i>addr</i>. If instruction and

data space are separated, request 4 writes a word into instruction space, and request 5 writes a word into data space. If instruction and data space are not separated, either request 4 or request 5 may be used with equal results. On success, the value written into the address space of the child is returned to the parent. These two requests fail if *addr* is not the start address of a word. On failure `-1` is returned to the parent process and the parent's `errno` is set to `EIO`.

- 6 With this request, a few entries in the child's user area can be written. *data* gives the value that is to be written and *addr* is the location of the entry. The few entries that can be written are the general registers and the condition codes of the Processor Status Word.
- 7 This request causes the child to resume execution. If the *data* argument is 0, all pending signals including the one that caused the child to stop are canceled before it resumes execution. If the *data* argument is a valid signal number, the child resumes execution as if it had incurred that signal, and any other pending signals are canceled. The *addr* argument must be equal to 1 for this request. On success, the value of *data* is returned to the parent. This request fails if *data* is not 0 or a valid signal number, in which case `-1` is returned to the parent process and the parent's `errno` is set to `EIO`.
- 8 This request causes the child to terminate with the same consequences as `exit(2)`.
- 9 This request sets the trace bit in the Processor Status Word of the child and then executes the same steps as listed above for request 7. The trace bit causes an interrupt on completion of one machine instruction. This effectively allows single stepping of the child.

To forestall possible fraud, `ptrace()` inhibits the set-user-ID facility on subsequent calls to one of the `exec` family of functions (see `exec(2)`). If a traced process calls one of the `exec` functions, it stops before executing the first instruction of the new image showing signal `SIGTRAP`.

ERRORS

The `ptrace()` function will fail if:

- | | |
|--------------|---|
| EIO | The <i>request</i> argument is an illegal number. |
| EPERM | The effective user of the calling process is not super-user. |
| ESRCH | The <i>pid</i> argument identifies a child that does not exist or has not executed a <code>ptrace()</code> call with request 0. |

SEE ALSO

`exec(2)`, `exit(2)`, `wait(2)`, `signal(3C)`, `signal(5)`

NAME	putmsg, putpmsg – send a message on a stream
SYNOPSIS	<pre>#include <stropts.h> int putmsg(int <i>fildev</i>, const struct strbuf * <i>ctlptr</i>, const struct strbuf * <i>dataptr</i>, int <i>flags</i>); int putpmsg(int <i>fildev</i>, const struct strbuf * <i>ctlptr</i>, const struct strbuf * <i>dataptr</i>, int <i>band</i>, int <i>flags</i>);</pre>
DESCRIPTION	<p>The putmsg() function creates a message 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.</p> <p>The putpmsg() function does the same thing as putmsg() , but provides the user the ability to send messages in different priority bands. Except where noted, all information pertaining to putmsg() also pertains to putpmsg() .</p> <p>The <i>fildev</i> argument specifies a file descriptor referencing an open stream. The <i>ctlptr</i> and <i>dataptr</i> arguments each point to a <i>strbuf</i> structure, which contains the following members:</p> <pre>int maxlen; /* not used here */ int len; /* length of data */ void *buf; /* ptr to buffer */</pre> <p>The <i>ctlptr</i> argument points to the structure describing the control part, if any, to be included in the message. The <i>buf</i> member in the <i>strbuf</i> structure points to the buffer where the control information resides, and the <i>len</i> member indicates the number of bytes to be sent. The <i>maxlen</i> member is not used in putmsg() (see getmsg(2)). In a similar manner, <i>dataptr</i> specifies the data, if any, to be included in the message. The <i>flags</i> argument indicates what type of message should be sent and is described later.</p> <p>To send the data part of a message, <i>dataptr</i> must not be <code>NULL</code> , and the <i>len</i> member of <i>dataptr</i> must have a value of 0 or greater. To send the control part of a message, the corresponding values must be set for <i>ctlptr</i> . No data (control) part is sent if either <i>dataptr</i> (<i>ctlptr</i>) is <code>NULL</code> or the <i>len</i> member of <i>dataptr</i> (<i>ctlptr</i>) is negative.</p> <p>For putmsg() , if a control part is specified, and <i>flags</i> is set to <code>RS_HIPRI</code> , a high priority message is sent. If no control part is specified, and <i>flags</i> is set to <code>RS_HIPRI</code> , putmsg() fails and sets <code>errno</code> to <code>EINVAL</code> . If <i>flags</i> is set to 0, a normal (non-priority) message is sent. If no control part and no data part are specified, and <i>flags</i> is set to 0, no message is sent, and 0 is returned.</p>

The stream head guarantees that the control part of a message generated by **putmsg()** is at least 64 bytes in length.

For **putpmsg()**, the flags are different. The *flags* argument is a bitmask with the following mutually-exclusive flags defined: `MSG_HIPRI` and `MSG_BAND`. If *flags* is set to 0, **putpmsg()** fails and sets `errno` to `EINVAL`. If a control part is specified and *flags* is set to `MSG_HIPRI` and *band* is set to 0, a high-priority message is sent. If *flags* is set to `MSG_HIPRI` and either no control part is specified or *band* is set to a non-zero value, **putpmsg()** fails and sets `errno` to `EINVAL`. If *flags* is set to `MSG_BAND`, then a message is sent in the priority band specified by *band*. If a control part and data part are not specified and *flags* is set to `MSG_BAND`, no message is sent and 0 is returned.

Normally, **putmsg()** will block if the stream write queue is full due to internal flow control conditions. For high-priority messages, **putmsg()** does not block on this condition. For other messages, **putmsg()** does not block when the write queue is full and `O_NDELAY` or `O_NONBLOCK` is set. Instead, it fails and sets `errno` to `EAGAIN`.

The **putmsg()** or **putpmsg()** function 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` or `O_NONBLOCK` has been specified. No partial message is sent.

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned and `errno` is set to indicate the error.

ERRORS

The **putmsg()** and **putpmsg()** functions will fail if:

EAGAIN	A non-priority message was specified, the <code>O_NDELAY</code> or <code>O_NONBLOCK</code> flag is set and the stream write queue is full due to internal flow control conditions.
EBADF	The <i>fildev</i> argument is not a valid file descriptor open for writing.
EFAULT	The <i>ctlptr</i> or <i>dataptr</i> argument points to an illegal address.
EINTR	A signal was caught during the execution of the putmsg() function.
EINVAL	An undefined value was specified in <i>flags</i> ; <i>flags</i> is set to <code>RS_HIPRI</code> and no control part was supplied; or the stream referenced by <i>fildev</i> is linked below a multiplexor.
ENOSR	Buffers could not be allocated for the message that was to be created due to insufficient STREAMS memory resources.

- ENOSTR** The *fildev* argument is not associated with a STREAM.
- ENXIO** A hangup condition was generated downstream for the specified stream, or the other end of the pipe is closed.
- EPIPE or EIO** The *fildev* argument refers to a STREAMS-based pipe and the other end of the pipe is closed. A SIGPIPE signal is generated for the calling process. This error condition occurs only with SUS-compliant applications. See **standards(5)** .
- 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 stream module. This value is also returned if the control part of the message is larger than the maximum configured size of the control part of a message, or if the data part of a message is larger than the maximum configured size of the data part of a message.

In addition, **putmsg()** and **putpmsg()** will fail if the STREAM head had processed an asynchronous error before the call. In this case, the value of `errno` does not reflect the result of **putmsg()** or **putpmsg()** but reflects the prior error.

The **putpmsg()** function will fail if:

- EINVAL** The *flags* argument is set to MSG_HIPRI and *band* is non-zero.

SEE ALSO

intro(2) , **getmsg(2)** , **poll(2)** , **read(2)** , **write(2)** , **standards(5)**

STREAMS Programming Guide

NAME	read, readv, pread – read from file
SYNOPSIS	<pre>#include <unistd.h> ssize_t read(int fildes, void * buf, size_t nbyte); ssize_t pread(int fildes, void * buf, size_t nbyte, off_t offset); #include <sys/uio.h> ssize_t readv(int fildes, const struct iovec * iov, int iovcnt);</pre>
DESCRIPTION	<p>The read() function attempts to read <i>nbyte</i> bytes from the file associated with the open file descriptor, <i>fildes</i> , into the buffer pointed to by <i>buf</i> .</p> <p>If <i>nbyte</i> is 0, read() will return 0 and have no other results.</p> <p>On files that support seeking (for example, a regular file), the read() starts at a position in the file given by the file offset associated with <i>fildes</i> . The file offset is incremented by the number of bytes actually read.</p> <p>Files that do not support seeking (for example, terminals) always read from the current position. The value of a file offset associated with such a file is undefined.</p> <p>If <i>fildes</i> refers to a socket, read() is equivalent to recv(3N) with no flags set.</p> <p>No data transfer will occur past the current end-of-file. If the starting position is at or after the end-of-file, 0 will be returned. If the file refers to a device special file, the result of subsequent read() requests is implementation-dependent.</p> <p>If the value of <i>nbyte</i> is greater than <code>SSIZE_MAX</code> , the result is implementation-dependent.</p> <p>When attempting to read from a regular file with mandatory file/record locking set (see chmod(2)), and there is a write lock owned by another process on the segment of the file to be read:</p> <ul style="list-style-type: none"> ■ If <code>O_NDELAY</code> or <code>O_NONBLOCK</code> is set, read() returns -1 and sets <code>errno</code> to <code>EAGAIN</code> . ■ If <code>O_NDELAY</code> and <code>O_NONBLOCK</code> are clear, read() sleeps until the blocking record lock is removed. <p>When attempting to read from an empty pipe (or FIFO):</p> <ul style="list-style-type: none"> ■ If no process has the pipe open for writing, read() returns 0 to indicate end-of-file. ■ If some process has the pipe open for writing and <code>O_NDELAY</code> is set, read() returns 0 .

- If some process has the pipe open for writing and `O_NONBLOCK` is set, **read()** returns `-1` and sets `errno` to `EAGAIN` .
- If `O_NDELAY` and `O_NONBLOCK` are clear, **read()** blocks until data is written to the pipe or the pipe is closed by all processes that had opened the pipe for writing.

When attempting to read a file associated with a terminal that has no data currently available:

- If `O_NDELAY` is set, **read()** returns `0` .
- If `O_NONBLOCK` is set, **read()** returns `-1` and sets `errno` to `EAGAIN` .
- If `O_NDELAY` and `O_NONBLOCK` are clear, **read()** blocks until data become available.

When attempting to read a file associated with a socket or a stream that is not a pipe, a FIFO, or a terminal, and the file has no data currently available:

- If `O_NDELAY` or `O_NONBLOCK` is set, **read()** returns `-1` and sets `errno` to `EAGAIN` .
- If `O_NDELAY` and `O_NONBLOCK` are clear, **read()** blocks until data becomes available.

The **read()** function reads data previously written to a file. If any portion of a regular file prior to the end-of-file has not been written, **read()** returns bytes with value `0`. For example, `lseek(2)` allows the file offset to be set beyond the end of existing data in the file. If data is later written at this point, subsequent reads in the gap between the previous end of data and the newly written data will return bytes with value `0` until data is written into the gap.

For regular files, no data transfer will occur past the offset maximum established in the open file description associated with *files* .

Upon successful completion, where *nbyte* is greater than `0`, **read()** will mark for update the `st_atime` field of the file, and return the number of bytes read. This number will never be greater than *nbyte* . The value returned may be less than *nbyte* if the number of bytes left in the file is less than *nbyte* , if the **read()** request was interrupted by a signal, or if the file is a pipe or FIFO or special file and has fewer than *nbyte* bytes immediately available for reading. For example, a **read()** from a file associated with a terminal may return one typed line of data.

If a **read()** is interrupted by a signal before it reads any data, it will return `-1` with `errno` set to `EINTR` .

If a **read()** is interrupted by a signal after it has successfully read some data, it will return the number of bytes read.

A **read()** from a STREAMS file can read data 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, and can be tested with the `I_GRDOPT ioctl()`. In byte-stream mode, **read()** retrieves 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()** retrieves data until as many bytes as were requested are transferred, or until a message boundary is reached. If **read()** does not retrieve all the data in a message, the remaining data is left on the STREAM, and can be retrieved by the next **read()** 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()** returns is discarded, and is not available for a subsequent **read()**, **readv()** or **getmsg(2)** call.

How **read()** handles zero-byte STREAMS messages is determined by the current read mode setting. In byte-stream mode, **read()** accepts data until it has read *nbyte* bytes, or until there is no more data to read, or until a zero-byte message block is encountered. The **read()** function then returns 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(2)**. In message-nondiscard mode or message-discard mode, a zero-byte message returns 0 and the message is removed from the STREAM. When a zero-byte message is read as the first message on a STREAM, the message is removed from the STREAM and 0 is returned, regardless of the read mode.

A **read()** from a STREAMS file returns the data in the message at the front of the STREAM head read queue, regardless of the priority band of the message.

By default, STREAMS are in control-normal mode, in which a **read()** from a STREAMS file can only process messages that contain a data part but do not contain a control part. The **read()** fails if a message containing a control part is encountered at the STREAM head. This default action can be changed by placing the STREAM in either control-data mode or control-discard mode with the `I_SRDOPT ioctl()` command. In control-data mode, **read()** converts any control part to data and passes it to the application before passing any data part originally present in the same message. In control-discard mode, **read()** discards message control parts but returns to the process any data part in the message.

In addition, **read()** and **readv()** will fail if the STREAM head had processed an asynchronous error before the call. In this case, the value of `errno` does not reflect the result of **read()** or **readv()** but reflects the prior error. If a hangup occurs on the STREAM being read, **read()** continues to operate normally until the STREAM head read queue is empty. Thereafter, it returns 0.

readv() The **readv()** function is equivalent to **read()**, but places the input data into the *iovcnt* buffers specified by the members of the *iov* array: *iov 0*, *iov 1*, ..., *iov [iovcnt - 1]*. The *iovcnt* argument is valid if greater than 0 and less than or equal to `IOV_MAX`.

The `iovec` structure contains the following members:

```
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. The **readv()** function always fills an area completely before proceeding to the next.

Upon successful completion, **readv()** marks for update the `st_atime` field of the file.

pread() The **pread()** function performs the same action as **read()**, except that it reads from a given position in the file without changing the file pointer. The first three arguments to **pread()** are the same as **read()** with the addition of a fourth argument *offset* for the desired position inside the file. **pread()** will read up to the maximum offset value that can be represented in an `off_t` for regular files. An attempt to perform a **pread()** on a file that is incapable of seeking results in an error.

RETURN VALUES

Upon successful completion, **read()** and **readv()** return a non-negative integer indicating the number of bytes actually read. Otherwise, the functions return `-1` and set `errno` to indicate the error.

ERRORS

The **read()**, **readv()**, and **pread()** functions will fail if:

- EAGAIN** Mandatory file/record locking was set, `O_NDELAY` or `O_NONBLOCK` was set, and there was a blocking record lock; total amount of system memory available when reading using raw I/O is temporarily insufficient; no data is waiting to be read on a file associated with a tty device and `O_NONBLOCK` was set; or no message is waiting to be read on a stream and `O_NDELAY` or `O_NONBLOCK` was set.
- EBADF** The *fdes* argument is not a valid file descriptor open for reading.
- EBADMSG** Message waiting to be read on a stream is not a data message.
- EDEADLK** The read was going to go to sleep and cause a deadlock to occur.

EFAULT	The <i>buf</i> argument points to an illegal address.
EINTR	A signal was caught during the read operation and no data was transferred.
EINVAL	An attempt was made to read from a stream linked to a multiplexor.
EIO	A physical I/O error has occurred, or the process is in a background process group and is attempting to read from its controlling terminal, and either the process is ignoring or blocking the <code>SIGTTIN</code> signal or the process group of the process is orphaned.
EISDIR	The <i>fildev</i> argument refers to a directory on a file system type that does not support read operations on directories.
ENOLCK	The system record lock table was full, so the <code>read()</code> or <code>readv()</code> could not go to sleep until the blocking record lock was removed.
ENOLINK	The <i>fildev</i> argument is on a remote machine and the link to that machine is no longer active.
ENXIO	The device associated with <i>fildev</i> is a block special or character special file and the value of the file pointer is out of range.
The <code>read()</code> and <code>readv()</code> functions will fail if:	
EOVERFLOW	The file is a regular file, <i>nbyte</i> is greater than 0, the starting position is before the end-of-file, and the starting position is greater than or equal to the offset maximum established in the open file description associated with <i>fildev</i> .
The <code>readv()</code> function may fail if:	
EFAULT	The <i>iov</i> argument points outside the allocated address space.
EINVAL	The <i>iovcnt</i> argument was less than or equal to 0, or greater than or equal to <code>{IOV_MAX}</code> . (See <code>intro(2)</code> for a definition of <code>{IOV_MAX}</code>).
EINVAL	The sum of the <i>iov_len</i> values in the <i>iov</i> array overflowed an int.
The <code>pread()</code> function will fail and the file pointer remain unchanged if:	
ESPIPE	The <i>fildev</i> argument is associated with a pipe or FIFO.

USAGE

The **pread()** function has a transitional interface for 64-bit file offsets. See **lf64(5)**.

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	read() is Async-Signal-Safe

SEE ALSO

intro(2), **chmod(2)**, **creat(2)**, **dup(2)**, **fcntl(2)**, **getmsg(2)**, **ioctl(2)**, **lseek(2)**, **open(2)**, **pipe(2)**, **recv(3N)**, **attributes(5)**, **lf64(5)**, **streamio(7I)**, **termio(7I)**

NAME	readlink – read the contents of a symbolic link
SYNOPSIS	<pre>#include <unistd.h> int readlink(const char *path, char *buf, size_t bufsiz);</pre>
DESCRIPTION	The readlink() function places the contents of the symbolic link referred to by <i>path</i> in the buffer <i>buf</i> which has size <i>bufsiz</i> . If the number of bytes in the symbolic link is less than <i>bufsiz</i> , the contents of the remainder of <i>buf</i> are unspecified.
RETURN VALUES	Upon successful completion, readlink() returns the count of bytes placed in the buffer. Otherwise, it returns <code>-1</code> , leaves the buffer unchanged, and sets <code>errno</code> to indicate the error.
ERRORS	<p>The readlink() function will fail if:</p> <p>EACCES Search permission is denied for a component of the path prefix of <i>path</i>.</p> <p>EFAULT <i>path</i> or <i>buf</i> points to an illegal address.</p> <p>EINVAL The <i>path</i> argument names a file that is not a symbolic link.</p> <p>EIO An I/O error occurred while reading from the file system.</p> <p>ENOENT A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.</p> <p>ELOOP Too many symbolic links were encountered in resolving <i>path</i>.</p> <p>ENAMETOOLONG The length of <i>path</i> exceeds <code>PATH_MAX</code>, or a pathname component is longer than <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.</p> <p>ENOTDIR A component of the path prefix is not a directory.</p> <p>ENOSYS The file system does not support symbolic links.</p> <p>The readlink() function may fail if:</p> <p>EACCES Read permission is denied for the directory.</p> <p>ENAMETOOLONG Path name resolution of a symbolic link produced an intermediate result whose length exceeds <code>PATH_MAX</code>.</p>

USAGE Portable applications should not assume that the returned contents of the symbolic link are null-terminated.

SEE ALSO `stat(2)`, `symlink(2)`

NAME	rename – change the name of a file
SYNOPSIS	<pre>#include <stdio.h> int rename(const char *old, const char *new);</pre>
DESCRIPTION	<p>The rename() function changes the name of a file. The <i>old</i> argument points to the pathname of the file to be renamed. The <i>new</i> argument points to the new pathname of the file.</p> <p>If <i>old</i> and <i>new</i> both refer to the same existing file, the rename() function returns successfully and performs no other action.</p> <p>If <i>old</i> points to the pathname of a file that is not a directory, <i>new</i> must not point to the pathname of a directory. If the link named by <i>new</i> exists, it will be removed and <i>old</i> will be renamed to <i>new</i>. In this case, a link named <i>new</i> must remain visible to other processes throughout the renaming operation and will refer to either the file referred to by <i>new</i> or the file referred to as <i>old</i> before the operation began.</p> <p>If <i>old</i> points to the pathname of a directory, <i>new</i> must not point to the pathname of a file that is not a directory. If the directory named by <i>new</i> exists, it will be removed and <i>old</i> will be renamed to <i>new</i>. In this case, a link named <i>new</i> will exist throughout the renaming operation and will refer to either the file referred to by <i>new</i> or the file referred to as <i>old</i> before the operation began. Thus, if <i>new</i> names an existing directory, it must be an empty directory.</p> <p>The <i>new</i> pathname must not contain a path prefix that names <i>old</i>. Write access permission is required for both the directory containing <i>old</i> and the directory containing <i>new</i>. If <i>old</i> points to the pathname of a directory, write access permission is required for the directory named by <i>old</i>, and, if it exists, the directory named by <i>new</i>.</p> <p>If the directory containing <i>old</i> has the sticky bit set, at least one of the following conditions listed below must be true:</p> <ul style="list-style-type: none">■ the user must own <i>old</i>■ the user must own the directory containing <i>old</i>■ <i>old</i> must be writable by the user■ the user must be a privileged user <p>If <i>new</i> exists, and the directory containing <i>new</i> is writable and has the sticky bit set, at least one of the following conditions must be true:</p> <ul style="list-style-type: none">■ the user must own <i>new</i>■ the user must own the directory containing <i>new</i>

- *new* must be writable by the user
- the user must be a privileged user

If the link named by *new* exists, the file's link count becomes zero when it is removed, and no process has the file open, then the space occupied by the file will be freed and the file will no longer be accessible. If one or more processes have the file open when the last link is removed, the link will be removed before **rename()** returns, but the removal of the file contents will be postponed until all references to the file have been closed.

Upon successful completion, the **rename()** function will mark for update the `st_ctime` and `st_mtime` fields of the parent directory of each file.

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned and `errno` is set to indicate an error.

ERRORS

The **rename()** function will fail if:

EACCES	A component of either path prefix denies search permission; one of the directories containing <i>old</i> and <i>new</i> denies write permissions; or write permission is denied by a directory pointed to by <i>old</i> or <i>new</i> .
EBUSY	The <i>new</i> argument is a directory and the mount point for a mounted file system.
EDQUOT	The directory where the new name entry is being placed cannot be extended because the user's quota of disk blocks on that file system has been exhausted.
EEXIST	The link named by <i>new</i> is a directory containing entries other than <code>.</code> (the directory itself) and <code>..</code> (the parent directory).
EINVAL	The <i>new</i> argument directory pathname contains a path prefix that names the <i>old</i> directory.
EISDIR	The <i>new</i> argument points to a directory but <i>old</i> points to a file that is not a directory.
ELOOP	Too many symbolic links were encountered in translating the pathname.
ENAMETOOLONG	The length of <i>old</i> or <i>new</i> exceeds <code>PATH_MAX</code> , or a pathname component is longer than <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.

EMLINK	The file named by <i>old</i> is a directory, and the link count of the parent directory of <i>new</i> would exceed <code>LINK_MAX</code> .
ENOENT	The link named by <i>old</i> does not exist, or either <i>old</i> or <i>new</i> points to an empty string.
ENOSPC	The directory that would contain <i>new</i> cannot be extended.
ENOTDIR	A component of either path prefix is not a directory, or <i>old</i> names a directory and <i>new</i> names a nondirectory file.
EROFS	The requested operation requires writing in a directory on a read-only file system.
EXDEV	The links named by <i>old</i> and <i>new</i> are on different file systems.
EIO	An I/O error occurred while making or updating a directory entry.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

`chmod(2)`, `link(2)`, `unlink(2)`, `attributes(5)`

NOTES

The system can deadlock if there is a loop in the file system graph. Such a loop can occur if there is an entry in directory *a*, *a/name1*, that is a hard link to directory *b*, and an entry in directory *b*, *b/name2*, that is a hard link to directory *a*. When such a loop exists and two separate processes attempt to rename *a/name1* to *b/name2* and *b/name2* to *a/name1*, the system may deadlock attempting to lock both directories for modification. Use symbolic links instead of hard links for directories.

NAME	resolvepath – resolve all symbolic links of a path name
SYNOPSIS	<pre>#include <unistd.h> int resolvepath(const char *path, char *buf, size_t bufsiz);</pre>
DESCRIPTION	<p>The resolvepath() function fully resolves all symbolic links in the path name <i>path</i> into a resulting path name free of symbolic links and places the resulting path name in the buffer <i>buf</i> which has size <i>bufsiz</i>. The resulting path name names the same file or directory as the original path name. All “.” components are eliminated and every non-leading “..” component is eliminated together with its preceding directory component. If leading “..” components reach to the root directory, they are replaced by “/”. If the number of bytes in the resulting path name is less than <i>bufsiz</i>, the contents of the remainder of <i>buf</i> are unspecified.</p>
RETURN VALUES	<p>Upon successful completion, resolvepath() returns the count of bytes placed in the buffer. Otherwise, it returns <code>-1</code>, leaves the buffer unchanged, and sets <code>errno</code> to indicate the error.</p>
ERRORS	<p>The resolvepath() function will fail if:</p> <p>EACCES Search permission is denied for a component of the path prefix of <i>path</i> or for a path prefix component resulting from the resolution of a symbolic link.</p> <p>EFAULT The <i>path</i> or <i>buf</i> argument points to an illegal address.</p> <p>EIO An I/O error occurred while reading from the file system.</p> <p>ENOENT The <i>path</i> argument is an empty string or a component of <i>path</i> or a path name component produced by resolving a symbolic link does not name an existing file.</p> <p>ELOOP Too many symbolic links were encountered in resolving <i>path</i>.</p> <p>ENAMETOOLONG The length of <i>path</i> exceeds <code>PATH_MAX</code>, or a path name component is longer than <code>NAME_MAX</code>. Path name resolution of a symbolic link produced an intermediate result whose length exceeds <code>PATH_MAX</code> or a component whose length exceeds <code>NAME_MAX</code>.</p>

ENOTDIR

A component of the path prefix of *path* or of a path prefix component resulting from the resolution of a symbolic link is not a directory.

USAGE

No more than `PATH_MAX` bytes will be placed in the buffer. Applications should not assume that the returned contents of the buffer are null-terminated.

SEE ALSO

`readlink(2)`, `realpath(3C)`

NAME	rmdir – remove a directory
SYNOPSIS	<pre>#include <unistd.h> int rmdir(const char *path);</pre>
DESCRIPTION	<p>The rmdir() function removes the directory named by the path name pointed to by <i>path</i>. The directory must not have any entries other than “.” and “..”.</p> <p>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.</p> <p>Upon successful completion rmdir() marks for update the <i>st_ctime</i> and <i>st_mtime</i> fields of the parent directory.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, <i>errno</i> is set to indicate the error, and the named directory is not changed.
ERRORS	<p>The rmdir() function will fail if:</p> <p>EACCES Search permission is denied for a component of the path prefix; write permission is denied on the directory containing the directory to be removed; the parent directory has the <i>S_ISVTX</i> variable set and is not owned by the user; the directory is not owned by the user and is not writable by the user; or the user is not a super-user.</p> <p>EBUSY The directory to be removed is the mount point for a mounted file system.</p> <p>EEXIST The directory contains entries other than those for “.” and “..”.</p> <p>EFAULT The <i>path</i> argument points to an illegal address.</p> <p>EINVAL The directory to be removed is the current directory, or the final component of <i>path</i> is “.”.</p> <p>EIO An I/O error occurred while accessing the file system.</p> <p>ELOOP Too many symbolic links were encountered in translating <i>path</i>.</p>

ENAMETOOLONG	The length of the <i>path</i> argument exceeds <code>PATH_MAX</code> , or the length of a <i>path</i> component exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.
ENOENT	The named directory does not exist or is the null pathname.
ENOLINK	The <i>path</i> argument points to a remote machine, and the connection to that machine is no longer active.
ENOTDIR	A component of the path prefix is not a directory.
EROFS	The directory entry to be removed is part of a read-only file system.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

`mkdir(1)`, `rm(1)`, `mkdir(2)`, `attributes(5)`

NAME	semctl – semaphore control operations
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ipc.h> #include <sys/sem.h> int semctl(int semid, int semnum, int cmd, ...);</pre>
DESCRIPTION	<p>The semctl() function provides a variety of semaphore control operations as specified by <i>cmd</i>. The fourth argument is optional, depending upon the operation requested. If required, it is of type <code>union semun</code>, which must be explicitly declared by the application program.</p> <pre>union semun { int val; struct semid_ds *buf; ushort_t *array; } arg ;</pre> <p>The permission required for a semaphore operation is given as <i>{token}</i>, where <i>token</i> is the type of permission needed. The types of permission are interpreted as follows:</p> <pre>00400 READ by user 00200 ALTER by user 00040 READ by group 00020 ALTER by group 00004 READ by others 00002 ALTER by others</pre> <p>See the Semaphore Operation Permissions subsection of the DEFINITIONS section of intro(2) for more information. The following semaphore operations as specified by <i>cmd</i> are executed with respect to the semaphore specified by <i>semid</i> and <i>semnum</i>.</p> <pre>GETVAL Return the value of semval (see intro(2)). {READ} SETVAL Set the value of semval to <i>arg.val</i>. {ALTER} When this command is successfully executed, the <i>semadj</i> value corresponding to the specified semaphore in all processes is cleared. GETPID Return the value of (int) sempid. {READ} GETNCNT Return the value of semncnt. {READ}</pre>

GETZCNT Return the value of `semzcnt`. {READ}
 The following operations return and set, respectively, every `semval` in the set of semaphores.

GETALL Place `semvals` into 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 operations 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 /* access permission bits only */
```

This command can be executed only 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 command 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`.

RETURN VALUES

Upon successful completion, the value returned depends on `cmd` as follows:

GETVAL the value of `semval`

GETPID the value of `(int) sempid`

GETNCNT the value of `semncnt`

GETZCNT the value of `semzcnt`

All other successful completions return 0; otherwise, -1 is returned and `errno` is set to indicate the error.

ERRORS	The semctl() function will fail if:
EACCES	Operation permission is denied to the calling process (see intro(2)).
EINVAL	The <i>semid</i> argument is not a valid semaphore identifier; the <i>semnum</i> argument is less than 0 or greater than <code>sem_nsems - 1</code> ; or the <i>cmd</i> argument is not a valid command or is <code>IPC_SET</code> and <code>sem_perm.uid</code> or <code>sem_perm.gid</code> is not valid.
EPERM	The <i>cmd</i> argument is equal to <code>IPC_RMID</code> or <code>IPC_SET</code> and the effective user of the calling process is not super-user, or <i>cmd</i> is equal to the value of <code>sem_perm.cuid</code> or <code>sem_perm.uid</code> in the data structure associated with <i>semid</i> .
E_OVERFLOW	The <i>cmd</i> argument is <code>IPC_STAT</code> and <i>uid</i> or <i>gid</i> is too large to be stored in the structure pointed to by <i>arg.buf</i> .
ERANGE	The <i>cmd</i> argument is <code>SETVAL</code> or <code>SETALL</code> and the value to which <i>semval</i> is to be set is greater than the system imposed maximum.

SEE ALSO `ipcs(1)`, `intro(2)`, `semget(2)`, `semop(2)`

NAME	semget – get set of semaphores
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ipc.h> #include <sys/sem.h> int semget(key_t key, int nsems, int semflg);</pre>
DESCRIPTION	<p>The semget() function returns the semaphore identifier associated with <i>key</i>.</p> <p>A semaphore identifier and associated data structure and set containing <i>nsems</i> semaphores (see intro(2)) are created for <i>key</i> if one of the following is true:</p> <ul style="list-style-type: none"> ■ <i>key</i> is equal to <code>IPC_PRIVATE</code>. ■ <i>key</i> does not already have a semaphore identifier associated with it, and $(semflg \& IPC_CREAT)$ is true. <p>On creation, the data structure associated with the new semaphore identifier is initialized as follows:</p> <ul style="list-style-type: none"> ■ <code>sem_perm.cuid</code>, <code>sem_perm.uid</code>, <code>sem_perm.cgid</code>, and <code>sem_perm.gid</code> are set equal to the effective user ID and effective group ID, respectively, of the calling process. ■ The access permission bits of <code>sem_perm.mode</code> are set equal to the access permission bits of <i>semflg</i>. ■ <code>sem_nsems</code> is set equal to the value of <i>nsems</i>. ■ <code>sem_otime</code> is set equal to 0 and <code>sem_ctime</code> is set equal to the current time.
RETURN VALUES	Upon successful completion, a non-negative integer representing a semaphore identifier is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The semget() function will fail if:</p> <p>EACCES A semaphore identifier exists for <i>key</i>, but operation permission (see intro(2)) as specified by the low-order 9 bits of <i>semflg</i> would not be granted.</p> <p>EEXIST A semaphore identifier exists for <i>key</i> but both $(semflg \& IPC_CREAT)$ and $(semflg \& IPC_EXCL)$ are both true.</p> <p>EINVAL The <i>nsems</i> argument is either less than or equal to 0 or greater than the system-imposed limit; or a semaphore identifier exists for <i>key</i>, but the number of semaphores in the set associated with it is less than <i>nsems</i> and <i>nsems</i> is not equal to 0.</p>

- 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 semaphores or semaphore identifiers system-wide would be exceeded.

SEE ALSO `ipcrm(1)`, `ipcs(1)`, `intro(2)`, `semctl(2)`, `semop(2)`, `ftok(3C)`

NAME	semop – semaphore operations
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ipc.h> #include <sys/sem.h> int semop(int semid, struct sembuf *sops, size_t nsops);</pre>
DESCRIPTION	<p>The semop() function is used to perform atomically an array of semaphore operations on the set of semaphores associated with the semaphore identifier specified by <i>semid</i>. The <i>sops</i> argument is a pointer to the array of semaphore-operation structures. The <i>nsops</i> argument is the number of such structures in the array.</p> <p>Each <i>sembuf</i> structure contains the following members:</p> <pre>short sem_num; /* semaphore number */ short sem_op; /* semaphore operation */ short sem_flg; /* operation flags */</pre> <p>Each semaphore operation specified by <i>sem_op</i> is performed on the corresponding semaphore specified by <i>semid</i> and <i>sem_num</i>. The permission required for a semaphore operation is given as <i>{token}</i>, where <i>token</i> is the type of permission needed. The types of permission are interpreted as follows:</p> <pre>00400 READ by user 00200 ALTER by user 00040 READ by group 00020 ALTER by group 00004 READ by others 00002 ALTER by others</pre> <p>See the <i>Semaphore Operation Permissions</i> section of intro(2) for more information.</p> <p>The <i>sem_op</i> member specifies one of three semaphore operations:</p> <ol style="list-style-type: none"> 1. The <i>sem_op</i> member is a negative integer; {ALTER} <ul style="list-style-type: none"> ■ If <i>semval</i> (see intro(2)) is greater than or equal to the absolute value of <i>sem_op</i>, the absolute value of <i>sem_op</i> is subtracted from <i>semval</i>. Also, if (<i>sem_flg</i>&SEM_UNDO) is true, the absolute value of <i>sem_op</i> is

added to the calling process's `semadj` value (see `exit(2)`) for the specified semaphore.

- If `semval` is less than the absolute value of `sem_op` and (`sem_flg&IPC_NOWAIT`) is true, **semop()** returns immediately.
 - If `semval` is less than the absolute value of `sem_op` and (`sem_flg&IPC_NOWAIT`) is false, **semop()** increments the `semncnt` associated with the specified semaphore and suspends execution of the calling process until one of the following conditions occur:
 - The value of `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 to `EIDRM` and `-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(3C)`.
2. The `sem_op` member is a positive integer; {ALTER}
- 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.
3. The `sem_op` member is 0; {READ}
- If `semval` is 0, **semop()** returns immediately.
 - If `semval` is not equal to 0 and (`sem_flg&IPC_NOWAIT`) is true, **semop()** returns immediately.
 - If `semval` is not equal to 0 and (`sem_flg&IPC_NOWAIT`) is false, **semop()** increments the `semzcnt` associated with the specified semaphore and suspends execution of the calling process until one of the following occurs:
 - The value of `semval` becomes 0, 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 to `EIDRM` and `-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(3C)`.

Upon successful completion, the value of `sempid` for each semaphore specified in the array pointed to by `sops` is set to the process ID of the calling process.

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned and `errno` is set to indicate the error.

ERRORS

The `semop()` function will fail if:

- E2BIG** The `nsops` argument is greater than the system-imposed maximum.
- EACCES** Operation permission is denied to the calling process (see `intro(2)`).
- EAGAIN** The operation would result in suspension of the calling process but (`sem_flg&IPC_NOWAIT`) is true.
- EFAULT** The `sops` argument points to an illegal address.
- EFBIG** The value of `sem_num` is less than 0 or greater than or equal to the number of semaphores in the set associated with `semid`.
- EIDRM** A `semid` was removed from the system.
- EINTR** A signal was received.
- EINVAL** The `semid` argument is not a valid semaphore identifier, or the number of individual semaphores for which the calling process requests a `SEM_UNDO` would exceed the limit.
- ENOSPC** The limit on the number of individual processes requesting an `SEM_UNDO` would be exceeded.
- ERANGE** An operation would cause a `semval` or a `semadj` value to overflow the system-imposed limit.

SEE ALSO

`ipcs(1)`, `intro(2)`, `exec(2)`, `exit(2)`, `fork(2)`, `semctl(2)`, `semget(2)`

NAME	setpgid – set process group ID
SYNOPSIS	<pre>#include <sys/types.h> #include <unistd.h> int setpgid(pid_t pid, pid_t pgid);</pre>
DESCRIPTION	<p>The setpgid() function sets the process group ID of the process with ID <i>pid</i> to <i>pgid</i>.</p> <p>If <i>pgid</i> is equal to <i>pid</i>, the process becomes a process group leader. See intro(2) for more information on session leaders and process group leaders.</p> <p>If <i>pgid</i> is not equal to <i>pid</i>, the process becomes a member of an existing process group.</p> <p>If <i>pid</i> is equal to 0, the process ID of the calling process is used. If <i>pgid</i> is equal to 0, the process specified by <i>pid</i> becomes a process group leader.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and errno is set to indicate the error.
ERRORS	<p>The setpgid() function will fail if:</p> <p>EACCES The <i>pid</i> argument matches the process ID of a child process of the calling process and the child process has successfully executed one of the <i>exec</i> family of functions (see exec(2)).</p> <p>EINVAL The <i>pgid</i> argument is less than (<code>pid_t</code>) 0 or greater than or equal to <code>PID_MAX</code>, or the calling process has a controlling terminal that does not support job control.</p> <p>EPERM The process indicated by the <i>pid</i> argument is a session leader.</p> <p>EPERM The <i>pid</i> argument 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.</p> <p>EPERM The <i>pgid</i> argument does not match the process ID of the process indicated by the <i>pid</i> argument, and there is no process with a process group ID that matches <i>pgid</i> in the same session as the calling process.</p> <p>ESRCH The <i>pid</i> argument does not match the process ID of the calling process or of a child process of the calling process.</p>
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

`intro(2)`, `exec(2)`, `exit(2)`, `fork(2)`, `getpid(2)`, `getsid(2)`,
`attributes(5)`

NAME	setpgrp – set process group ID
SYNOPSIS	<pre>#include <sys/types.h> #include <unistd.h> pid_t setpgrp(void);</pre>
DESCRIPTION	If the calling process is not already a session leader, the setpgrp() function makes it one by setting its process group ID and session ID to the value of its process ID, and releases its controlling terminal. See intro(2) for more information on process group IDs and session leaders.
RETURN VALUES	The setpgrp() function returns the value of the new process group ID.
SEE ALSO	intro(2) , exec(2) , fork(2) , getpid(2) , getsid(2) , kill(2) , signal(3C)
NOTES	The setpgrp() function will be phased out in favor of the setsid(2) function.

NAME	setregid – set real and effective group IDs
SYNOPSIS	<pre>#include <unistd.h> int setregid(gid_t rgid, gid_t egid);</pre>
DESCRIPTION	<p>The setregid() function is used to set the real and effective group IDs of the calling process. If <i>rgid</i> is -1, the real group ID is not changed; if <i>egid</i> is -1, the effective group ID is not changed. The real and effective group IDs may be set to different values in the same call.</p> <p>If the effective user ID of the calling process is super-user, the real group ID and the effective group ID can be set to any legal value.</p> <p>If the effective user ID of the calling process is not super-user, either the real group ID can be set to the saved set-group-ID from execve(2), or the effective group ID can either be set to the saved set-group-ID or the real group ID.</p> <p>In either case, if the real group ID is being changed (that is, if <i>rgid</i> is not -1), or the effective group ID is being changed to a value not equal to the real group ID, the saved set-group-ID is set equal to the new effective group ID.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, <i>errno</i> is set to indicate the error, and neither of the group IDs will be changed.
ERRORS	<p>The setregid() function will fail if:</p> <p>EINVAL The value of <i>rgid</i> or <i>egid</i> is less than 0 or greater than <code>UID_MAX</code> (defined in <code><limits.h></code>).</p> <p>EPERM The calling process's effective UID is not the super-user and a change other than changing the real group ID to the saved set-group-ID or changing the effective group ID to the real group ID or the saved group ID, was specified.</p>
USAGE	If a set-group-ID process sets its effective group ID to its real group ID, it can still set its effective group ID back to the saved set-group-ID.
SEE ALSO	execve(2) , getgid(2) , setreuid(2) , setuid(2)

NAME	setreuid – set real and effective user IDs
SYNOPSIS	<pre>#include <unistd.h> int setreuid(uid_t ruid, uid_t euid);</pre>
DESCRIPTION	<p>The setreuid() function is used to set the real and effective user IDs of the calling process. If <i>ruid</i> is -1, the real user ID is not changed; if <i>euid</i> 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.</p> <p>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.</p> <p>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() (see exec(2)) or the real user ID.</p> <p>In either case, if the real user ID is being changed (that is, if <i>ruid</i> 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.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, <i>errno</i> is set to indicate the error, and neither of the user IDs will be changed.
ERRORS	<p>The setreuid() function will fail if:</p> <p>EINVAL The value of <i>ruid</i> or <i>euid</i> is less than 0 or greater than <code>UID_MAX</code> (defined in <code><limits.h></code>).</p> <p>EPERM The calling process's 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.</p>
USAGE	If a set-user-ID 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.
SEE ALSO	exec(2) , getuid(2) , setregid(2) , setuid(2)

NAME	setsid – create session and set process group ID				
SYNOPSIS	<pre>#include <sys/types.h> #include <unistd.h> pid_t setsid(void);</pre>				
DESCRIPTION	The setsid() function creates a new session, if the calling process is not a process group leader. Upon return the calling process will be the session leader of this new session, will be the process group leader of a new process group, and will have no controlling terminal. The process group ID of the calling process will be set equal to the process ID of the calling process. The calling process will be the only process in the new process group and the only process in the new session.				
RETURN VALUES	Upon successful completion, setsid() returns the value of the process group ID of the calling process. Otherwise it returns <code>(pid_t)-1</code> and sets <code>errno</code> to indicate the error.				
ERRORS	<p>The setsid() function will fail if:</p> <p>EPERM The calling process is already a process group leader, or the process group ID of a process other than the calling process matches the process ID of the calling process.</p>				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	getsid(2) , setpgid(2) , setpgrp(2) , attributes(5)				
WARNINGS	A call to setsid() by a process that is a process group leader will fail. A process can become a process group leader by being the last member of a pipeline started by a job control shell. Thus, a process that expects to be part of a pipeline, and that calls setsid() , should always first fork; the parent should exit and the child should call setsid() . This will ensure that the calling process will work reliably when started by both job control shells and non-job control shells.				

NAME	setuid, setegid, seteuid, setgid – set user and group IDs
SYNOPSIS	<pre>#include <sys/types.h> #include <unistd.h> int setuid(uid_t uid); int setegid(gid_t egid); int seteuid(uid_t euid); int setgid(gid_t gid);</pre>
DESCRIPTION	<p>The setuid() function sets the real user ID , effective user ID , and saved user ID of the calling process. The setgid() function sets the real group ID , effective group ID , and saved group ID of the calling process. The setegid() and seteuid() functions set the effective group and user ID s respectively for the calling process. See intro(2) for more information on real, effective, and saved user and group ID s.</p> <p>At login time, the real user ID , effective user ID , and saved user ID of the login process are set to the login ID of the user responsible for the creation of the process. The same is true for the real, effective, and saved group ID s; they are set to the group ID of the user responsible for the creation of the process.</p> <p>When a process calls one of the exec family of functions (see exec(2))to execute a file (program), the user and/or group identifiers associated with the process can change. If the file executed is a set-user- ID file, the effective and saved user ID s of the process are set to the owner of the file executed. If the file executed is a set-group- ID file, the effective and saved group ID s of the process are set to the group of the file executed. If the file executed is not a set-user- ID or set-group- ID file, the effective user ID , saved user ID , effective group ID , and saved group ID are not changed.</p> <p>If the effective user ID of the process calling setuid() is the super-user, the real, effective, and saved user IDs are set to the <i>uid</i> argument.</p> <p>If the effective user ID of the calling process is not the super-user, but <i>uid</i> is either the real user ID or the saved user ID of the calling process, the effective user ID is set to <i>uid</i> .</p> <p>If the effective user ID of the process calling setgid() is the super-user, the real, effective, and saved group ID s are set to the <i>gid</i> argument.</p> <p>If the effective user ID of the calling process is not the super-user, but <i>gid</i> is either the real group ID or the saved group ID of the calling process, the effective group ID is set to <i>gid</i> .</p>

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned and `errno` is set to indicate the error.

ERRORS

The `setuid()` and `setgid()` functions will fail if:

EINVAL

The value of `uid` or `gid` is out of range.

EPERM

For `setuid()` and `seteuid()` the effective user of the calling process is not super-user, and the `uid` argument does not match either the real or saved user ID s. For `setgid()` and `setegid()` the effective user of the calling process is not the super-user, and the `gid` argument does not match either the real or saved group ID s.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	<code>setuid()</code> and <code>setgid()</code> and Async-Signal-Safe

SEE ALSO

`intro(2)`, `exec(2)`, `getgroups(2)`, `getuid(2)`, `attributes(5)`, `stat(5)`

NAME	shmctl – shared memory control operations
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ipc.h> #include <sys/shm.h> int shmctl(int shmid, int cmd, struct shmids *buf);</pre>
DESCRIPTION	<p>The shmctl() function provides a variety of shared memory control operations as specified by <i>cmd</i>. The permission required for a shared memory control operation is given as <i>{token}</i>, where <i>token</i> is the type of permission needed. The types of permission are interpreted as follows:</p> <pre>00400 READ by user 00200 WRITE by user 00040 READ by group 00020 WRITE by group 00004 READ by others 00002 WRITE by others</pre> <p>See the <i>Shared Memory Operation Permissions</i> section of intro(2) for more information.</p> <p>The following operations require the specified tokens:</p> <p>IPC_STAT Place the current value of each member of the data structure associated with <i>shmid</i> into the structure pointed to by <i>buf</i>. The contents of this structure are defined in intro(2). {READ}</p> <p>IPC_SET Set the value of the following members of the data structure associated with <i>shmid</i> to the corresponding value found in the structure pointed to by <i>buf</i>:</p> <pre>shm_perm.uid shm_perm.gid shm_perm.mode /* access permission bits only */</pre> <p>This command can be executed only by a process that has an effective user ID equal to that of super-user, or to the value of <i>shm_perm.cuid</i> or <i>shm_perm.uid</i> in the data structure associated with <i>shmid</i>.</p> <p>IPC_RMID Remove the shared memory identifier specified by <i>shmid</i> from the system and destroy the shared memory segment and data structure associated with it. This command can be executed only by a process that has an effective user ID equal to that of super-user, or to the value of <i>shm_perm.cuid</i> or <i>shm_perm.uid</i> in the data structure associated with <i>shmid</i>.</p>

SHM_LOCK Lock the shared memory segment specified by *shmid* in memory. This command can be executed only by a process that has an effective user ID equal to super-user.

SHM_UNLOCK Unlock the shared memory segment specified by *shmid*. This command can be executed only by a process that has an effective user ID equal to super-user.

Shared memory segments must be explicitly removed after the last reference to them has been removed.

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The **shmctl()** function will fail if:

EACCES The *cmd* argument is equal to `IPC_STAT` and `{READ}` operation permission is denied to the calling process.

EFAULT The *buf* argument points to an illegal address.

EINVAL The *shmid* argument is not a valid shared memory identifier; or the *cmd* argument is not a valid command or is `IPC_SET` and *shm_perm.uid* or *shm_perm.gid* is not valid.

ENOMEM The *cmd* argument is equal to `SHM_LOCK` and there is not enough memory.

EOVERFLOW The *cmd* argument is `IPC_STAT` and *uid* or *gid* is too large to be stored in the structure pointed to by *buf*.

EPERM The *cmd* argument is equal to `IPC_RMID` or `IPC_SET` and the effective user ID of the calling process is not super-user and it is not equal to the value of *shm_perm.cuid* or *shm_perm.uid* in the data structure associated with *shmid*.

EPERM The *cmd* argument is equal to `SHM_LOCK` or `SHM_UNLOCK` and the effective user ID of the calling process is not equal to that of super-user.

SEE ALSO

ipcs(1), *intro(2)*, *shmget(2)*, *shmop(2)*

NAME	shmget – get shared memory segment identifier
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/ipc.h> #include <sys/shm.h> int shmget(key_t key, size_t size, int shmflg);</pre>
DESCRIPTION	<p>The shmget() function returns the shared memory identifier associated with <i>key</i>.</p> <p>A shared memory identifier and associated data structure and shared memory segment of at least <i>size</i> bytes (see intro(2)) are created for <i>key</i> if one of the following are true:</p> <ul style="list-style-type: none"> ■ The <i>key</i> argument is equal to <code>IPC_PRIVATE</code>. ■ The <i>key</i> argument does not already have a shared memory identifier associated with it, and $(shmflg \& IPC_CREAT)$ is true. <p>Upon creation, the data structure associated with the new shared memory identifier is initialized as follows:</p> <ul style="list-style-type: none"> ■ The values of <code>shm_perm.cuid</code>, <code>shm_perm.uid</code>, <code>shm_perm.cgid</code>, and <code>shm_perm.gid</code> are set equal to the effective user ID and effective group ID, respectively, of the calling process. ■ The access permission bits of <code>shm_perm.mode</code> are set equal to the access permission bits of <i>shmflg</i>. <code>shm_segsz</code> is set equal to the value of <i>size</i>. ■ The values of <code>shm_lpid</code>, <code>shm_nattch</code>, <code>shm_atime</code>, and <code>shm_dtime</code> are set equal to 0. ■ The <code>shm_ctime</code> is set equal to the current time. <p>Shared memory segments must be explicitly removed after the last reference to them has been removed.</p>
RETURN VALUES	<p>Upon successful completion, a non-negative integer representing a shared memory identifier is returned. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.</p>
ERRORS	<p>The shmget() function will fail if:</p> <p>EACCES A shared memory identifier exists for <i>key</i> but operation permission (see intro(2)) as specified by the low-order 9 bits of <i>shmflg</i> would not be granted.</p> <p>EEXIST A shared memory identifier exists for <i>key</i> but both $(shmflg \& IPC_CREATE)$ and $(shmflg \& IPC_EXCL)$ are true.</p>

- EINVAL** The *size* argument is less than the system-imposed minimum or greater than the system-imposed maximum.
- EINVAL** A shared memory identifier exists for *key* but the size of the segment associated with it is less than *size* and *size* is not equal to 0.
- ENOENT** A shared memory identifier does not exist for *key* and (*shmflg*&IPC_CREATE) is false.
- ENOMEM** A shared memory identifier and associated shared memory segment are to be created but the amount of available 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 `intro(2)`, `shmctl(2)`, `shmop(2)`, `ftok(3C)`

NAME	shmop, shmat, shmdt – shared memory operations
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/shm.h> void * shmat(int <i>shmid</i>, const void * <i>shmaddr</i>, int <i>shmflg</i>);</pre>
Default	int shmdt (char * <i>shmaddr</i>);
Standard-conforming	int shmdt (const void * <i>shmaddr</i>);
DESCRIPTION	<p>The shmat() function attaches the shared memory segment associated with the shared memory identifier specified by <i>shmid</i> to the data segment of the calling process.</p> <p>The permission required for a shared memory control operation is given as { <i>token</i> }, where <i>token</i> is the type of permission needed. The types of permission are interpreted as follows:</p> <pre>00400 READ by user 00200 WRITE by user 00040 READ by group 00020 WRITE by group 00004 READ by others 00002 WRITE by others</pre> <p>See the <i>Shared Memory Operation Permissions</i> section of intro(2) for more information.</p> <p>When (<i>shmflg</i> & SHM_SHARE_MMU) is true, virtual memory resources in addition to shared memory itself are shared among processes that use the same shared memory.</p> <p>The shared memory segment is attached to the data segment of the calling process at the address specified based on one of the following criteria:</p> <ul style="list-style-type: none"> ■ If <i>shmaddr</i> is equal to (void *) 0 , the segment is attached to the first available address as selected by the system. ■ If <i>shmaddr</i> is equal to (void *) 0 and (<i>shmflg</i> & SHM_SHARE_MMU) is true, then the segment is attached to the first available suitably aligned address. When (<i>shmflg</i> & SHM_SHARE_MMU) is set, however, the permission given by shmget() determines whether the segment is attached for reading or reading and writing. ■ If <i>shmaddr</i> is not equal to (void *) 0 and (<i>shmflg</i> & SHM_RND) is true, the segment is attached to the address given by (<i>shmaddr</i> - (<i>shmaddr</i> modulus SHMLBA)).

- If *shmaddr* is not equal to `(void *) 0` and `(shmflg & SHM_RND)` is false, the segment is attached to the address given by *shmaddr*.
- The segment is attached for reading if `(shmflg & SHM_RDONLY)` is true {READ}, otherwise it is attached for reading and writing {READ/WRITE}.

The **shmdt()** function detaches from the calling process's data segment the shared memory segment located at the address specified by *shmaddr*. If the application is standard-conforming (see **standards(5)**), the *shmaddr* argument is of type `const void *`. Otherwise it is of type `char *`.

Shared memory segments must be explicitly removed after the last reference to them has been removed.

RETURN VALUES

Upon successful completion, **shmat()** returns the data segment start address of the attached shared memory segment; **shmdt()** returns 0. Otherwise, -1 is returned, the shared memory segment is not attached, and `errno` is set to indicate the error.

ERRORS

The **shmat()** function will fail if:

- | | |
|---------------|--|
| EACCES | Operation permission is denied to the calling process (see intro(2)). |
| EINVAL | The <i>shmid</i> argument is not a valid shared memory identifier. |
| EINVAL | The <i>shmaddr</i> argument is not equal to 0, and the value of <code>(shmaddr - (shmaddr modulus SHMLBA))</code> is an illegal address. |
| EINVAL | The <i>shmaddr</i> argument is not equal to 0, is an illegal address, and <code>(shmflg & SHM_RND)</code> is false. |
| EINVAL | The <i>shmaddr</i> argument is not equal to 0, is not properly aligned, and <code>(shmflg & SHM_SHARE_MMU)</code> is true. |
| EINVAL | SHM_SHARE_MMU is not supported in certain architectures. |
| EMFILE | The number of shared memory segments attached 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. |

The **shmdt()** function will fail if:

- | | |
|---------------|---|
| EINVAL | The <i>shmaddr</i> argument is not the data segment start address of a shared memory segment. |
|---------------|---|

SEE ALSO | `intro(2)`, `exec(2)`, `exit(2)`, `fork(2)`, `shmctl(2)`, `shmget(2)`,
`standards(5)`

NAME	sigaction – detailed signal management
SYNOPSIS	<pre>#include <signal.h> int sigaction(int sig, const struct sigaction *act, struct sigaction *oact);</pre>
DESCRIPTION	<p>The sigaction() function allows the calling process to examine or specify the action to be taken on delivery of a specific signal. (See signal(5) for an explanation of general signal concepts.)</p> <p>The <i>sig</i> argument specifies the signal and can be assigned any of the signals specified in signal(5) except SIGKILL and SIGSTOP. In a multithreaded process, <i>sig</i> cannot be SIGWAITING, SIGCANCEL, or SIGLWP.</p> <p>If the argument <i>act</i> is not NULL, it points to a structure specifying the new action to be taken when delivering <i>sig</i>. If the argument <i>oact</i> is not NULL, it points to a structure where the action previously associated with <i>sig</i> is to be stored on return from sigaction().</p> <p>The sigaction structure includes the following members:</p> <pre>void (*sa_handler)(); void (*sa_sigaction)(int, siginfo_t *, void *); sigset_t sa_mask; int sa_flags;</pre> <p>The <i>sa_handler</i> member identifies the action to be associated with the specified signal, if the SA_SIGINFO flag (see below) is cleared in the <i>sa_flags</i> field of the sigaction structure. It may take any of the values specified in signal(5) or that of a user specified signal handler. If the SA_SIGINFO flag is set in the <i>sa_flags</i> field, the <i>sa_sigaction</i> field specifies a signal-catching function.</p> <p>The <i>sa_mask</i> member specifies a set of signals to be blocked while the signal handler is active. On entry to the signal handler, that set of signals is added to the set of signals already being blocked when the signal is delivered. In addition, the signal that caused the handler to be executed will also be blocked, unless the SA_NODEFER flag has been specified. SIGSTOP and SIGKILL cannot be blocked (the system silently enforces this restriction).</p> <p>The <i>sa_flags</i> member specifies a set of flags used to modify the delivery of the signal. It is formed by a logical OR of any of the following values:</p> <p>SA_ONSTACK If set and the signal is caught, and if the LWP that is chosen to process a delivered signal has an alternate signal stack declared with sigaltstack(2), then it will process the signal on that stack. Otherwise, the signal is delivered on the LWP main stack. Unbound threads (see thr_create(3T)) may not have alternate signal stacks.</p>

SA_RESETHAND	If set and the signal is caught, the disposition of the signal is reset to SIG_DFL and the signal will not be blocked on entry to the signal handler (SIGILL, SIGTRAP, and SIGPWR cannot be automatically reset when delivered; the system silently enforces this restriction).
SA_NODEFER	If set and the signal is caught, the signal will not be automatically blocked by the kernel while it is being caught.
SA_RESTART	If set and the signal is caught, functions that are interrupted by the execution of this signal's handler are transparently restarted by the system, namely <code>fcntl(2)</code> , <code>ioctl(2)</code> , <code>wait(2)</code> , <code>waitid(2)</code> , and the following functions on slow devices like terminals: <code>getmsg()</code> and <code>getpmsg()</code> (see <code>getmsg(2)</code>); <code>putmsg()</code> and <code>putpmsg()</code> (see <code>putmsg(2)</code>); <code>pread()</code> , <code>read()</code> , and <code>readv()</code> (see <code>read(2)</code>); <code>pwrite()</code> , <code>write()</code> , and <code>writv()</code> (see <code>write(2)</code>); <code>recv()</code> , <code>recvfrom()</code> , and <code>recvmsg()</code> (see <code>recv(3N)</code>); and <code>send()</code> , <code>sendto()</code> , and <code>sendmsg()</code> (see <code>recv(3N)</code>). Otherwise, the function returns an EINTR error.
SA_SIGINFO	If cleared and the signal is caught, <code>sig</code> is passed as the only argument to the signal-catching function. If set and the signal is caught, two additional arguments are passed to the signal-catching function. If the second argument is not equal to NULL, it points to a <code>siginfo_t</code> structure containing the reason why the signal was generated (see <code>siginfo(5)</code>); the third argument points to a <code>ucontext_t</code> structure containing the receiving process's context when the signal was delivered (see <code>ucontext(5)</code>).
SA_NOCLDWAIT	If set and <code>sig</code> equals SIGCHLD, the system will not create zombie processes when children of the calling process exit. If the calling process subsequently issues a <code>wait(2)</code> , it blocks until all of the calling process's child processes terminate, and then returns <code>-1</code> with <code>errno</code> set to ECHILD.
SA_NOCLDSTOP	If set and <code>sig</code> equals SIGCHLD, SIGCHLD will not be sent to the calling process when its child processes stop or continue.
SA_WAITSIG	If set and <code>sig</code> equals SIGWAITING, enables generation of SIGWAITING signals. Reserved for use by the threads library.

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, `-1` is returned, `errno` is set to indicate the error, and no new signal handler is installed.

ERRORS

The **sigaction()** function will fail if:

EINVAL The value of the *sig* argument is not a valid signal number or is equal to SIGKILL or SIGSTOP. In addition, if in a multithreaded process, it is equal to SIGWAITING, SIGCANCEL, or SIGLWP.

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

kill(1), **intro(2)**, **exit(2)**, **fcntl(2)**, **getmsg(2)**, **ioctl(2)**, **kill(2)**, **pause(2)**, **putmsg(2)**, **read(2)**, **sigaltstack(2)**, **sigprocmask(2)**, **sigsend(2)**, **sigsuspend(2)**, **wait(2)**, **waitid(2)**, **write(2)**, **recv(3N)**, **recv(3N)**, **signal(3C)**, **sigsetops(3C)**, **thr_create(3T)**, **attributes(5)**, **siginfo(5)**, **signal(5)**, **ucontext(5)**

NOTES

The handler routine can be declared:

```
void handler (int sig, siginfo_t *sip, ucontext_t *uap);
```

The *sig* argument is the signal number. The *sip* argument is a pointer (to space on the stack) to a siginfo_t structure, which provides additional detail about the delivery of the signal. The *uap* argument is a pointer (again to space on the stack) to a ucontext_t structure (defined in <sys/ucontext.h>) which contains the context from before the signal. It is not recommended that *uap* be used by the handler to restore the context from before the signal delivery.

NAME	sigaltstack – set or get signal alternate stack context
SYNOPSIS	<pre>#include <signal.h> int sigaltstack(const stack_t *ss, stack_t *oss);</pre>
DESCRIPTION	<p>The sigaltstack() function allows an LWP to define an alternate stack area on which signals are to be processed. If <i>ss</i> is non-zero, it specifies a pointer to, and the size of a stack area on which to deliver signals, and tells the system whether the LWP is currently executing on that stack. When a signal's action indicates its handler should execute on the alternate signal stack (specified with a sigaction(2) call), the system checks to see if the LWP chosen to execute the signal handler is currently executing on that stack. If the LWP is not currently executing on the signal stack, the system arranges a switch to the alternate signal stack for the duration of the signal handler's execution.</p> <p>The <code>stack_t</code> structure includes the following members:</p> <pre>int *ss_sp long ss_size int ss_flags</pre> <p>If <i>ss</i> is not <code>NULL</code>, it points to a structure specifying the alternate signal stack that will take effect upon successful return from sigaltstack(). The <code>ss_sp</code> and <code>ss_size</code> members specify the new base and size of the stack, which is automatically adjusted for direction of growth and alignment. The <code>ss_flags</code> member specifies the new stack state and may be set to the following:</p> <p><code>SS_DISABLE</code> The stack is to be disabled and <code>ss_sp</code> and <code>ss_size</code> are ignored. If <code>SS_DISABLE</code> is not set, the stack will be enabled.</p> <p>If <i>oss</i> is not <code>NULL</code>, it points to a structure specifying the alternate signal stack that was in effect prior to the call to sigaltstack(). The <code>ss_sp</code> and <code>ss_size</code> members specify the base and size of that stack. The <code>ss_flags</code> member specifies the stack's state, and may contain the following values:</p> <p><code>SS_ONSTACK</code> The LWP is currently executing on the alternate signal stack. Attempts to modify the alternate signal stack while the LWP is executing on it will fail.</p> <p><code>SS_DISABLE</code> The alternate signal stack is currently disabled.</p>
RETURN VALUES	Upon successful completion, 0 is return. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The sigaltstack() function will fail if:</p> <p>EFAULT The <i>ss</i> or <i>oss</i> argument points to an illegal address.</p>

- EINVAL** The *ss* argument is not a null pointer, and the *ss_flags* member pointed to by *ss* contains flags other than *SS_DISABLE*.
- ENOMEM** The size of the alternate stack area is less than *MINSIGSTKSZ*.
- EPERM** An attempt was made to modify an active stack.

SEE ALSO

`getcontext(2)`, `sigaction(2)`, `ucontext(5)`

NOTES

The value *SIGSTKSZ* is defined to be the number of bytes that would be used to cover the usual case when allocating an alternate stack area. The value *MINSIGSTKSZ* is defined to be the minimum stack size for a signal handler. In computing an alternate stack size, a program should add that amount to its stack requirements to allow for the operating system overhead.

The following code fragment is typically used to allocate an alternate stack:

```
if ((sigstk.ss_sp = (char *)malloc(SIGSTKSZ)) == NULL)
    /* error return */;

sigstk.ss_size = SIGSTKSZ;
sigstk.ss_flags = 0;
if (sigaltstack(&sigstk, (stack_t *)0) < 0)
    perror("sigaltstack");
```

NAME	_signotifywait, _lwp_sigredirect – deliver process signals to specific LWPs
SYNOPSIS	<pre>#include <sys/lwp.h> int _signotifywait(void); int _lwp_sigredirect(lwpid_t target_lwp, int signo);</pre>
DESCRIPTION	<p>In a multithreaded process, signals that are generated for a process are delivered to one of the threads that does not have that signal masked. If all of the application threads are masking that signal, its delivery waits until one of them un.masks it.</p> <p>The disposition of the each thread's signal mask is unknown to the kernel when it generates signals for the process. The _signotifywait() and _lwp_sigredirect() functions provide a mechanism to direct instances of signals generated for the process to application-specified LWPs. Each process has a set of signals pending for the process, and for each LWP there is a set of signals pending for that LWP. If no signals are pending, these sets are empty.</p> <p>There is also a process-wide signal set, termed the <i>notification</i> set, manipulated by these functions. A signal generated for the process where the signal number is not in the notification set is called an <i>unnotified</i> signal.</p> <p>In a multithreaded program there is an <code>aslwp</code>, a special LWP endowed with powers to handle signals that are generated for a process. The _signotifywait() function is used to await signals generated for the process, and should be called only from the <code>aslwp</code>. In general, these functions are not to be called from the application-level.</p> <p>If there is a pending unnotified signal when _signotifywait() is called, that signal is selected and the call returns immediately. If there is not a signal pending, the call suspends the calling LWP until the generation of an unnotified signal; that signal then is selected and the function returns. In both cases, the selected signal number is set in the notification set and returned as the value of _signotifywait(). The signal remains pending for the process, and any associated <code>siginfo(5)</code> information remains queued at the process.</p> <p>The _lwp_sigredirect() function requests that a signal pending for the process be delivered to the LWP specified by <code>target_lwp</code>. If <code>target_lwp</code> is 0, the signal is discarded. It is an error if <code>signo</code> is not currently in the notification set of the process. The signal specified by <code>signo</code> is removed from pending for the process and is made pending for the <code>target_lwp</code>. If there is an associated <code>siginfo</code> information structure queued at the process, that <code>siginfo</code> is queued to the <code>target_lwp</code>.</p> <p>Whenever a signal is cleared from the set of signals pending for the process, the corresponding signal is cleared from the notification set. After a successful</p>

call to **_lwp_sigredirect()** , the signal *signo* is cleared from the notification set and from the set of signals pending for the process. If another instance of *signo* is queued for the process, the signal number is again set in the process pending mask, and if another LWP is blocked in a call to **_signotifywait()** , its wait for an unnotified signal will be satisfied. The effects described in this paragraph also apply when the signal *signo* is returned by a call to **sigtimedwait()** and *signo* was not pending for the calling LWP.

RETURN VALUES

The **_signotifywait()** function returns the signal number of the pending but hitherto unnotified signal. The **_lwp_sigredirect()** function returns 0 when successful. A non-zero value indicates an error.

ERRORS

No error conditions are specified for **_signotifywait()** .

If the following conditions occurs, **_lwp_sigredirect()** fails and return the corresponding value:

EINVAL The signal *signo* was not pending for the process, or *signo* was not in the notification set.

ESRCH The *target_lwp* cannot be found in the current process.

SEE ALSO

_lwp_create(2) , **_lwp_kill(2)** , **sigtimedwait(3R)** , **siginfo(5)** , **signal(5)**

NOTES

This mechanism for delivering signals to multithreaded processes is subject to change in future versions of Solaris. Any process with explicit knowledge of this mechanism may not be compatible from release to release.

NAME	sigpending – examine signals that are blocked and pending				
SYNOPSIS	#include <signal.h> int sigpending(sigset_t *set);				
DESCRIPTION	The sigpending() function retrieves those signals that have been sent to the calling process but are being blocked from delivery by the calling process's signal mask. The signals are stored in the space pointed to by the <i>set</i> argument.				
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <i>errno</i> is set to indicate the error.				
ERRORS	The sigpending() function will fail if: EFAULT The <i>set</i> argument points to an illegal address.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	sigaction(2) , sigprocmask(2) , sigsetops(3C) , attributes(5)				

NAME	sigprocmask – change and/or examine caller’s signal mask				
SYNOPSIS	<pre>#include <signal.h> int sigprocmask(int how, const sigset_t *set, sigset_t *oset);</pre>				
DESCRIPTION	<p>The sigprocmask() function is used to examine and/or change the caller’s signal mask. If the value is <code>SIG_BLOCK</code>, the set pointed to by the <i>set</i> argument is added to the current signal mask. If the value is <code>SIG_UNBLOCK</code>, the set pointed by the <i>set</i> argument is removed from the current signal mask. If the value is <code>SIG_SETMASK</code>, the current signal mask is replaced by the set pointed to by the <i>set</i> argument. If the <i>oset</i> argument is not <code>NULL</code>, the previous mask is stored in the space pointed to by <i>oset</i>. If the value of the <i>set</i> argument is <code>NULL</code>, the value <i>how</i> is not significant and the caller’s signal mask is unchanged; thus, the call can be used to inquire about currently blocked signals.</p> <p>If there are any pending unblocked signals after the call to sigprocmask(), at least one of those signals will be delivered before the call to sigprocmask() returns.</p> <p>It is not possible to block those signals that cannot be ignored this restriction is silently imposed by the system. See sigaction(2).</p> <p>If sigprocmask() fails, the caller’s signal mask is not changed.</p>				
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.				
ERRORS	<p>The sigprocmask() function will fail if:</p> <p>EFAULT The <i>set</i> or <i>oset</i> argument points to an illegal address.</p> <p>EINVAL The value of the <i>how</i> argument is not equal to one of the defined values.</p>				
ATTRIBUTES	<p>See attributes(5) for descriptions of the following attributes:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="text-align: center;">MT-Level</td> <td style="text-align: center;">Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	sigaction(2) , signal(3C) , sigsetops(3C) , thr_sigsetmask(3T) , attributes(5) , signal(5)				
NOTES	In a multithreaded program, the call to sigpromask() impacts only the calling thread’s signal mask and is therefore identical to a call to thr_sigsetmask(3T) .				

Signals that are generated synchronously should not be masked. If such a signal is blocked and delivered, the receiving process is killed.

NAME	sigsend, sigsendset – send a signal to a process or a group of processes
SYNOPSIS	<pre>#include <signal.h> int sigsend(idtype_t idtype, id_t id, int sig); int sigsendset(procset_t * psp, int sig);</pre>
DESCRIPTION	<p>The sigsend() function sends a signal to the process or group of processes specified by <i>id</i> and <i>idtype</i> . The signal to be sent is specified by <i>sig</i> and is either 0 or one of the values listed in signal(5) . If <i>sig</i> is 0 (the null signal), error checking is performed but no signal is actually sent. This value can be used to check the validity of <i>id</i> and <i>idtype</i> .</p> <p>The real or effective user ID of the sending process must match the real or saved user ID of the receiving process, unless the effective user ID of the sending process is super-user, or <i>sig</i> is SIGCONT and the sending process has the same session ID as the receiving process.</p> <p>If <i>idtype</i> is P_PID , <i>sig</i> is sent to the process with process ID <i>id</i> .</p> <p>If <i>idtype</i> is P_PGID , <i>sig</i> is sent to all process with process group ID <i>id</i> .</p> <p>If <i>idtype</i> is P_SID , <i>sig</i> is sent to all process with session ID <i>id</i> .</p> <p>If <i>idtype</i> is P_UID , <i>sig</i> is sent to any process with effective user ID <i>id</i> .</p> <p>If <i>idtype</i> is P_GID , <i>sig</i> is sent to any process with effective group ID <i>id</i> .</p> <p>If <i>idtype</i> is P_CID , <i>sig</i> is sent to any process with scheduler class ID <i>id</i> (see pricnt1(2)).</p> <p>If <i>idtype</i> is P_ALL , <i>sig</i> is sent to all processes and <i>id</i> is ignored.</p> <p>If <i>id</i> is P_MYID , the value of <i>id</i> is taken from the calling process.</p> <p>The process with a process ID of 0 is always excluded. The process with a process ID of 1 is excluded unless <i>idtype</i> is equal to P_PID .</p> <p>The sigsendset() function provides an alternate interface for sending signals to sets of processes. This function sends signals to the set of processes specified by <i>psp</i> . <i>psp</i> is a pointer to a structure of type <code>procset_t</code> , defined in <code><sys/procset.h></code> , which includes the following members:</p> <pre>idop_t p_op; idtype_t p_lidtype; id_t p_lid; idtype_t p_ridtype; id_t p_rid;</pre>

The `p_lidtype` and `p_lid` members specify the ID type and ID of one (“left”) set of processes; the `p_ridtype` and `p_rid` members specify the ID type and ID of a second (“right”) set of processes. ID types and IDs are specified just as for the `idtype` and `id` arguments to `sigsend()`. The `p_op` member specifies the operation to be performed on the two sets of processes to get the set of processes the function is to apply to. The valid values for `p_op` and the processes they specify are:

<code>POP_DIFF</code>	Set difference: processes in left set and not in right set.
<code>POP_AND</code>	Set intersection: processes in both left and right sets.
<code>POP_OR</code>	Set union: processes in either left or right set or both.
<code>POP_XOR</code>	Set exclusive-or: processes in left or right set but not in both.

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned and `errno` is set to indicate the error.

ERRORS

The `sigsend()` and `sigsendset()` functions will fail if:

EINVAL	The <code>sig</code> argument is not a valid signal number, or the <code>idtype</code> argument is not a valid idtype field.
EINVAL	The <code>sig</code> argument is <code>SIGKILL</code> , <code>idtype</code> is <code>P_PID</code> and <code>id</code> is 1 (<code>proc1</code>).
EPERM	The effective user of the calling process is not super-user and its real or effective user ID does not match the real or effective user ID of the receiving process, and the calling process is not sending <code>SIGCONT</code> to a process that shares the same session.
ESRCH	No process can be found corresponding to that specified by <code>id</code> and <code>idtype</code> .

The `sigsendset()` function will fail if:

EFAULT	The <code>psp</code> argument points to an illegal address.
---------------	---

SEE ALSO

`kill(1)`, `getpid(2)`, `kill(2)`, `prctl(2)`, `signal(3C)`, `signal(5)`

NAME	sigsuspend – install a signal mask and suspend caller until signal				
SYNOPSIS	<pre>#include <signal.h> int sigsuspend(const sigset_t *set);</pre>				
DESCRIPTION	<p>The sigsuspend() function replaces the caller's signal mask with the set of signals pointed to by the <i>set</i> argument and suspends the caller until delivery of a signal whose action is either to execute a signal catching function or to terminate the process.</p> <p>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. On return, the signal mask is restored to the set that existed before the call to sigsuspend(). See NOTES for the precise semantics of signal mask restoration in a multithreaded process.</p> <p>It is not possible to block those signals that cannot be ignored (see signal(5)); this restriction is silently imposed by the system.</p>				
RETURN VALUES	Since sigsuspend() suspends process execution indefinitely, there is no successful completion return value. On failure, it returns <code>-1</code> and sets <code>errno</code> to indicate the error.				
ERRORS	<p>The sigsuspend() function will fail if:</p> <p>EFAULT The <i>set</i> argument points to an illegal address.</p> <p>EINTR A signal was caught by the calling process and control was returned from the signal catching function.</p>				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	sigaction(2) , sigprocmask(2) , sigwait(2) , signal(3C) , sigsetops(3C) , attributes(5) , signal(5)				
NOTES	In a multithreaded application, the sigwait(2) , function should be used instead of sigsuspend() . Should sigsuspend() be used, however, its semantics of signal mask restoration are slightly different from those for a single-threaded process on return from the signal catching function, the signal mask is restored to the set that existed before the call to sigsuspend() . This action raises the following implications:				

- If a thread specifies two signals in the mask to **sigsuspend()**, both signals could interrupt its call to **sigsuspend()** simultaneously. In the traditional program that does not use threads, a call to **sigsuspend()** with two signals in the mask always returns with only one signal delivered. The other signal remains pending if masked earlier, unlike the MT case.
- While a thread is executing the signal handler that interrupted its call to **sigsuspend()**, its signal mask is the one passed to **sigsuspend()**. It does not get restored to the previous mask until it returns from all the signal handlers that interrupted **sigsuspend()**.

NAME	sigwait – wait until a signal is posted
SYNOPSIS	
Default	#include <signal.h> int sigwait(sigset_t *set);
POSIX	cc [<i>flag ...</i>] <i>file ...</i> -D_POSIX_PTHREAD_SEMANTICS [<i>library...</i>] #include <signal.h> int sigwait(const sigset_t *set, int *sig);
DESCRIPTION	<p>The sigwait() function selects a signal in <i>set</i> that is pending on the calling thread (see thr_create(3T)) or LWP. If no signal in <i>set</i> is pending, then sigwait() blocks until a signal in <i>set</i> becomes pending. The selected signal is cleared from the set of signals pending on the calling thread or LWP and the number of the signal is returned, or in the POSIX version (see standards(5)) placed in <i>sig</i>. The selection of a signal in <i>set</i> is independent of the signal mask of the calling thread or LWP. This means a thread or LWP can synchronously wait for signals that are being blocked by the signal mask of the calling thread or LWP. To ensure that only the caller receives the signals defined in <i>set</i>, all threads should have signals in <i>set</i> masked including the calling thread.</p> <p>If sigwait() is called on an ignored signal, then the occurrence of the signal will be ignored, unless sigaction() changes the disposition. If more than one thread or LWP waits for the same signal, only one is unblocked when the signal arrives.</p>
RETURN VALUES	Upon successful completion, the default version of sigwait() returns a signal number; the POSIX version returns 0 and stores the received signal number at the location pointed to by <i>sig</i> . Otherwise, -1 is returned and <i>errno</i> is set to indicate an error.
ERRORS	<p>The sigwait() function will fail if:</p> <p>EINVAL The <i>set</i> argument contains an unsupported signal number.</p> <p>EFAULT The <i>set</i> argument points to an invalid address.</p>
EXAMPLES	<p>EXAMPLE 1 Creating a thread to handle receipt of a signal</p> <p>The following sample C code creates a thread to handle the receipt of a signal. More specifically, it catches the asynchronously generated signal, SIGINT.</p> <pre> /***** * * compile with -D_POSIX_PTHREAD_SEMANTICS switch; * required by sigwait() * *****/ </pre>

```

* sigint thread handles delivery of signal. uses sigwait() to wait
* for SIGINT signal.
*
*****/
#include <pthread.h>
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <unistd.h>
#include <signal.h>
#include <synch.h>

static void    *threadTwo(void *);
static void    *threadThree(void *);
static void    *sigint(void *);

sigset_t      signalSet;

void *
main(void)
{
    pthread_t    t;
    pthread_t    t2;
    pthread_t    t3;

    thr_setconcurrency(3);
    sigfillset ( &signalSet );
    /*
     * Block signals in initial thread. New threads will
     * inherit this signal mask.
     */
    pthread_sigmask ( SIG_BLOCK, &signalSet, NULL );

    printf("Creating threads\n");

    /* POSIX thread create arguments:
     * thr_id, attr, strt_func, arg
     */
    pthread_create(&t, NULL, sigint, NULL);
    pthread_create(&t2, NULL, threadTwo, NULL);
    pthread_create(&t3, NULL, threadThree, NULL);

    printf("#####\n");
    printf("press CTRL-C to deliver SIGINT to sigint thread\n");
    printf("#####\n");

    thr_exit((void *)0);
}

static void *
threadTwo(void *arg)
{
    printf("hello world, from threadTwo [tid: %d]\n",
           pthread_self());
    printf("threadTwo [tid: %d] is now complete and exiting\n",
           pthread_self());
    thr_exit((void *)0);
}

```

```

static void *
threadThree(void *arg)
{
    printf("hello world, from threadThree [tid: %d]\n",
           pthread_self());
    printf("threadThree [tid: %d] is now complete and exiting\n",
           pthread_self());
    thr_exit((void *)0);
}

void *
sigint(void *arg)
{
    int    sig;
    int    err;

    printf("thread sigint [tid: %d] awaiting SIGINT\n",
           pthread_self());

    /* use POSIX sigwait() -- 2 args
     * signal set, signum
     */
    err = sigwait ( &signalSet, &sig );

    /* test for SIGINT; could catch other signals */
    if (err || sig != SIGINT)
        abort();

    printf("\nSIGINT signal %d caught by sigint thread [tid: %d]\n",
           sig, pthread_self());
    thr_exit((void *)0);
}

```

SEE ALSO

sigaction(2), **sigpending(2)**, **sigprocmask(2)**, **sigsuspend(2)**, **thr_create(3T)**, **thr_sigsetmask(3T)**, **signal(5)**, **standards(5)**

NOTES

The **sigwait()** function cannot be used to wait for signals that cannot be caught (see **sigaction(2)**). This restriction is silently imposed by the system.

Solaris 2.4 and earlier releases provided a **sigwait()** facility as specified in POSIX.1c Draft 6. The final POSIX.1c standard changed the interface as described above. Support for the Draft 6 interface is provided for compatibility only and may not be supported in future releases. New applications and libraries should use the POSIX standard interface.

In Solaris 2.4 and earlier releases, the call to **sigwait()** from a multithreaded process overrode the signal's ignore disposition; even if a signal's disposition was **SIG_IGN**, a call to **sigwait()** resulted in catching the signal, if generated. This is unspecified behavior from the standpoint of the POSIX 1003.1c specification.

In Solaris 2.5, the behavior of **sigwait()** was corrected, so that it does not override the signal's ignore disposition. This change can cause applications that rely on the old behavior to break. Applications should employ **sigwait()** as follows: Install a dummy signal handler, thereby changing the disposition from SIG_IGN to having a handler. Then, any calls to **sigwait()** for this signal would catch it upon generation.

NAME | `__sparc_utrap_install` - install a SPARC V9 user trap handler

SYNOPSIS | `#include <sys/utrap.h>`

`int __sparc_utrap_install(utrap_entry_t type, utrap_handler_t new_precise,
utrap_handler_t new_deferred, utrap_handler_t *old_precise, utrap_handler_t *old_deferred);`

DESCRIPTION | The `__sparc_utrap_install()` function establishes *new_precise* and *new_deferred* user trap handlers as the new values for the specified *type* and returns the existing user trap handler values in **old_precise* and **old_deferred* in a single atomic operation. A new handler address of `NULL` means no user handler of that type will be installed. A new handler address of `UTH_NOCHANGE` means that the user handler for that type should not be changed. An old handler pointer of `NULL` means that the user is not interested in the old handler address.

A *precise trap* is caused by a specific instruction and occurs before any program-visible state has been changed by this instruction. When a precise trap occurs, the program counter (PC) saved in the Trap Program Counter (TPC) register points to the instruction that induced the trap; all instructions prior to this trapping instruction have been executed. The next program counter (nPC) saved in the Trap Next Program Counter (TnPC) register points to the next instruction following the trapping instruction, which has not yet been executed. A *deferred trap* is also caused by a particular instruction, but unlike a precise trap, a deferred trap may occur after the program-visible state has been changed. See the *SPARC Architecture Manual, Version 9* for further information on precise and deferred traps.

The list that follows contains hardware traps and their corresponding user trap types. User trap types marked with a plus-sign (+) are required and must be provided by all ABI-conforming implementations. The others may not be present on every implementation; an attempt to install a user trap handler for those conditions will return `EINVAL`. User trap types marked with an asterisk (*) are implemented as precise traps only.

Trap Name	User Trap Type (utrap_entry_t)
<code>illegal_instruction</code>	<code>UT_ILLTRAP_INSTRUCTION +*</code> <code>UT_ILLEGAL_INSTRUCTION</code>
<code>fp_disabled</code>	<code>UT_FP_DISABLED +*</code>
<code>fp_exception_ieee_754</code>	<code>UT_FP_EXCEPTION_IEEE_754 +</code>
<code>fp_exception_other</code>	<code>UT_FP_EXCEPTION_OTHER</code>
<code>tag_overflow</code>	<code>UT_TAG_OVERFLOW +*</code>
<code>division_by_zero</code>	<code>UT_DIVISION_BY_ZERO +</code>
<code>mem_address_not_aligned</code>	<code>UT_MEM_ADDRESS_NOT_ALIGNED +</code>

Trap Name	User Trap Type (utrap_entry_t)
privileged_action	UT_PRIVILEGED_ACTION +
privileged_opcode	UT_PRIVILEGED_OPCODE
async_data_error	UT_ASYNC_DATA_ERROR
trap_instruction	UT_TRAP_INSTRUCTION_16 through UT_TRAP_INSTRUCTION_31 +*
instruction_access_exception instruction_access_MMU_miss instruction_access_error	UT_INSTRUCTION_EXCEPTION or UT_INSTRUCTION_PROTECTION or UT_INSTRUCTION_ERROR
data_access_exception data_access_MMU_miss data_access_error data_access_protection	UT_DATA_EXCEPTION or UT_DATA_PROTECTION or UT_DATA_ERROR

The following explanations are provided for those user trap types that are not self-explanatory.

UT_ILLTRAP_INSTRUCTION

This trap is raised by user execution of the ILLTRAP INSTRUCTION. It is always precise.

UT_ILLEGAL_INSTRUCTION

This trap will be raised by the execution of otherwise undefined opcodes. It is implementation-dependent as to what opcodes raise this trap; the ABI only specifies the interface. The trap may be precise or deferred.

UT_PRIVILEGED_OPCODE

All opcodes declared to be privileged in SPARC V9 will raise this trap. It is implementation-dependent whether other opcodes will raise it as well; the ABI only specifies the interface.

UT_DATA_EXCEPTION, UT_INSTRUCTION_EXCEPTION

No valid user mapping can be made to this address, for a data or instruction access, respectively.

UT_DATA_PROTECTION, UT_INSTRUCTION_PROTECTION

A valid mapping exists, and user privilege to it exists, but the type of access (read, write, or execute) is denied, for a data or instruction access, respectively.

UT_DATA_ERROR, UT_INSTRUCTION_ERROR

A valid mapping exists, and both user privilege and the type of access are allowed, but an unrecoverable error occurred in attempting the access, for a data or instruction access, respectively. %11 will contain either BUS_ADDRERR or BUS_OBJERR.

UT_FP_DISABLED

This trap is raised when an application issues a floating point instruction (including load or store) and the SPARC V9 Floating Point Registers State (FPRS) FEF bit is 0. If a user handler is installed for this trap, it will be given control. Otherwise the system will set FEF to one and retry the instruction. For all traps, the handler executes in a new register window, where the *in* registers are the *out* registers of the previous frame and have the value they contained at the time of the trap, similar to a normal subroutine call after the *save* instruction. The *global* registers (including the special registers %CCR, %asi, and %y) and the *floating-point* registers have their values from the time of the trap. The stack pointer register %sp plus the BIAS will point to a properly-aligned 128-byte register save area; if the handler needs scratch space, it should decrement the stack pointer to obtain it. If the handler needs access to the previous frame's *in* registers or *local* registers, it should execute a FLUSHW instruction, and then access them off of the frame pointer. If the handler calls an ABI-conforming function, it must set the %asi register to ASI_PRIMARY_NOFAULT before the call.

On entry to a precise user trap handler %16 contains the %pc and %17 contains the %npc at the time of the trap. To return from a handler and reexecute the trapped instruction, the handler would execute:

```
jmp1 %16, %g0 ! Trapped PC supplied to user trap handler
return %17    ! Trapped nPC supplied to user trap handler
```

To return from a handler and skip the trapped instruction, the handler would execute:

```
jmp1 %17, %g0 ! Trapped nPC supplied to user trap handler
return %17 + 4 ! Trapped nPC + 4
```

On entry to a deferred trap handler %o0 contains the address of the instruction that caused the trap and %o1 contains the actual instruction (right-justified, zero-extended), if the information is available. Otherwise %o0 contains the

value -1 and %o1 is undefined. Additional information may be made available for certain cases of deferred traps, as indicated in the following table.

Instructions	Additional Information
LD-type LDSTUB	%o2 contains the effective address (rs1 + rs2 simm13).
ST-type CAS SWAP	%o2 contains the effective address (rs1 + rs2 simm13).
Integer arithmetic	%o2 contains the rs1 value. %o3 contains the rs2 simm13 value. %o4 contains the contents of the %Y register.
Floating-point arithmetic	%o2 contains the address of rs1 value. %o3 contains the address of rs2 value.
Control-transfer	%o2 contains the target address (rs1 + rs2 simm13).
Asynchronous data errors	%o2 contains the address that caused the error. %o3 contains the effective ASI, if available, else -1.

To return from a deferred trap, the trap handler issues:

```
ta 68 !ST_RETURN_FROM_DEFERRED_TRAP
```

The following pseudo-code explains how the operating system dispatches traps:

```
if (precise trap) {
    if (precise_handler) {
        invoke(precise_handler);
        /* not reached */
    } else {
        convert_to_signal(precise_trap);
    }
} else if (deferred_trap) {
    invoke(deferred_handler);
    /* not reached */
} else {
    convert_to_signal(deferred_trap);
}
}
if (signal)
    send(signal);
```

User trap handlers must preserve all registers except the *locals* (%l0-7) and the *outs* (%o0-7), that is, %i0-7, %g1-7, %d0-d62, %asi, %fsr, %fprs, %ccr, and %y, except to the extent that modifying the registers is part of the desired functionality of the handler. For example, the handler for UT_FP_DISABLED may load floating-point registers.

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, a non-zero value is returned and `errno` is set to indicate the error.

ERRORS

The `__sparc_utrap_install()` function will fail if:

EINVAL The *type* argument is not a supported user trap type; the new user trap handler address is not word aligned; the old user trap handler address cannot be returned; or the user program is not a 64-bit executable.

EXAMPLES

EXAMPLE 1 A sample program using the `__sparc_utrap_install()` function.

The `__sparc_utrap_install()` function is normally used by user programs that wish to provide their own tailored exception handlers as a faster alternative to `signal(2)`, or to handle exceptions that are not directly supported by the `signal()` interface, such as `fp_disabled`.

```
extern void *fpdis_trap_handler();
utrap_handler_t new_precise = (utrap_handler_t)fpdis_trap_handler;
double d;
int err;
err = __sparc_utrap_install(UT_FP_DISABLED, new_precise,
    UTH_NOCHANGE, NULL, NULL);
if (err == EINVAL) {
    /* unexpected error, do something */
    exit (1);
}
d = 1.0e-300;
ENTRY(fpdis_trap_handler)
wr    %g0, FPRS_FEF, %fprs
jmpl  %l6, %g0
return %l7
SET_SIZE(fpdis_trap_handler)
```

This example turns on bit 2, FEF, in the Floating-Point Registers State (FPRS) Register, after a floating-point instruction causes an `fp_disabled` trap. (Note that this example simulates part of the default system behavior; programs do not need such a handler. The example is for illustrative purposes only.)

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO

signal(2), **attributes(5)**

SPARC Architecture Manual, Version 9

Manufacturer's processor chip user manuals

NOTES

The Exceptions and Interrupt Descriptions section of the SPARC V9 manual documents which hardware traps are mandatory or optional, and whether they can be implemented as precise or deferred traps, or both. The manufacturer's processor chip user manuals describe the details of the traps supported for the specific processor implementation.

NAME	stat, lstat, fstat – get file status
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/stat.h> int stat(const char * path, struct stat * buf); int lstat(const char * path, struct stat * buf); int fstat(int fildes, struct stat * buf);</pre>
DESCRIPTION	<p>The stat() function obtains information about the file pointed to by <i>path</i>. 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.</p> <p>The lstat() function obtains file attributes similar to stat(), except when the named file is a symbolic link; in that case lstat() returns information about the link, while stat() returns information about the file the link references.</p> <p>The fstat() function obtains information about an open file known by the file descriptor <i>fildes</i>, obtained from a successful open(2), creat(2), dup(2), fcntl(2), or pipe(2) function.</p> <p>The <i>buf</i> argument is a pointer to a stat structure into which information is placed concerning the file. A stat structure includes the following members:</p> <pre>mode_t st_mode; /* File mode (see mknod(2)) */ ino_t st_ino; /* Inode number */ dev_t st_dev; /* ID of device containing */ /* a directory entry for this file */ dev_t st_rdev; /* ID of device */ /* This entry is defined only for */ /* char special or block special files */ nlink_t st_nlink; /* Number of links */ uid_t st_uid; /* User ID of the file's owner */ gid_t st_gid; /* Group ID of the file's group */ off_t st_size; /* File size in bytes */ time_t st_atime; /* Time of last access */ time_t st_mtime; /* Time of last data modification */ time_t st_ctime; /* Time of last file status change */ /* Times measured in seconds since */ /* 00:00:00 UTC, Jan. 1, 1970 */ long st_blksize; /* Preferred I/O block size */ blkcnt_t st_blocks; /* Number of 512 byte blocks allocated*/</pre> <p>Descriptions of structure members are as follows:</p> <p>st_mode The mode of the file as described in mknod(2). In addition to the modes described in mknod(), the mode of a file may also be S_IFLNK if the file is a symbolic link. S_IFLNK may only be returned by lstat().</p>

st_ino	This field uniquely identifies the file in a given file system. The pair st_ino and st_dev uniquely identifies regular files.
st_dev	This field uniquely identifies the file system that contains the file. Its value may be used as input to the ustat() function to determine more information about this file system. No other meaning is associated with this value.
st_rdev	This field should be used only by administrative commands. It is valid only for block special or character special files and only has meaning on the system where the file was configured.
st_nlink	This field should be used only by administrative commands.
st_uid	The user ID of the file's owner.
st_gid	The group ID of the file's group.
st_size	For regular files, this is the address of the end of the file. For block special or character special, this is not defined. See also pipe(2) .
st_atime	Time when file data was last accessed. Changed by the following functions: creat() , mknod() , pipe() , utime(2) , and read(2) .
st_mtime	Time when data was last modified. Changed by the following functions: creat() , mknod() , pipe() , utime() , and write(2) .
st_ctime	Time when file status was last changed. Changed by the following functions: chmod() , chown() , creat() , link(2) , mknod() , pipe() , unlink(2) , utime() , and write() .
st_blksize	A hint as to the "best" unit size for I/O operations. This field is not defined for block special or character special files.
st_blocks	The total number of physical blocks of size 512 bytes actually allocated on disk. This field is not defined for block special or character special files.

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned and **errno** is set to indicate the error.

ERRORS

The **stat()** , **fstat()** , and **lstat()** functions will fail if:

EOVERFLOW The file size in bytes or the number of blocks allocated to the file or the file serial number cannot be represented correctly in the structure pointed to by *buf* .

The **stat()** and **lstat()** functions will fail if:

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

EFAULT The *buf* or *path* argument points to an illegal address.

EINTR A signal was caught during the execution of the **stat()** or **lstat()** function.

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

ENAMETOOLONG The length of the *path* argument exceeds `PATH_MAX` , or the length of a *path* component exceeds `NAME_MAX` while `_POSIX_NO_TRUNC` is in effect.

ENOENT The named file does not exist or is the null pathname.

ENOLINK The *path* argument points to a remote machine and the link to that machine is no longer active.

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

EOVERFLOW A component is too large to store in the structure pointed to by *buf* .

The **fstat()** function will fail if:

EBADF The *fildev* argument is not a valid open file descriptor.

EFAULT The *buf* argument points to an illegal address.

EINTR A signal was caught during the execution of the **fstat()** function.

ENOLINK The *fildev* argument points to a remote machine and the link to that machine is no longer active.

EOVERFLOW A component is too large to store in the structure pointed to by *buf* .

USAGE The `stat()`, `fstat()`, and `lstat()` functions have transitional interfaces for 64-bit file offsets. See `lf64(5)`.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	<code>stat()</code> and <code>fstat()</code> are Async-Signal-Safe

SEE ALSO `chmod(2)`, `chown(2)`, `creat(2)`, `link(2)`, `mknod(2)`, `pipe(2)`, `read(2)`, `time(2)`, `unlink(2)`, `utime(2)`, `write(2)`, `fattach(3C)`, `attributes(5)`, `lf64(5)`, `stat(5)`

NOTES If you use `chmod(2)` to change the file group owner permissions on a file with ACL entries, both the file group owner permissions and the ACL mask are changed to the new permissions. Be aware that the new ACL mask permissions may change the effective permissions for additional users and groups who have ACL entries on the file.

NAME statvfs, fstatvfs – get file system information

SYNOPSIS #include <sys/types.h>
#include <sys/statvfs.h>

```
int statvfs(const char * path, struct statvfs * buf);
```

```
int fstatvfs(int fildes, struct statvfs * buf);
```

DESCRIPTION

The **statvfs()** function returns a “generic superblock” describing a file system; it can be used to acquire information about mounted file systems. The *buf* argument is a pointer to a structure (described below) that is filled by the function.

The *path* argument should name a file that resides on that file system. The file system type is known to the operating system. Read, write, or execute permission for the named file is not required, but all directories listed in the path name leading to the file must be searchable.

The *statvfs* structure pointed to by *buf* includes the following members:

```
u_long      f_bsize;           /* preferred file system block size */
u_long      f_frsize;         /* fundamental filesystem block
                               (size if supported) */
fsblkcnt_t  f_blocks;         /* total # of blocks on file system
                               in units of f_frsize */
fsblkcnt_t  f_bfree;          /* total # of free blocks */
fsblkcnt_t  f_bavail;         /* # of free blocks avail to
                               non-super-user */
fsfilcnt_t  f_files;          /* total # of file nodes (inodes) */
fsfilcnt_t  f_ffree;          /* total # of free file nodes */
fsfilcnt_t  f_favail;         /* # of inodes avail to
                               non-super-user*/
u_long      f_fsid;           /* file system id (dev for now) */
char        f_basetype[FSTYPSZ];
/* target fs type name,
                               null-terminated */
u_long      f_flag;           /* bit mask of flags */
u_long      f_namemax;        /* maximum file name length */
char        f_fstr[32];       /* file system specific string */
u_long      f_filler[16];     /* reserved for future expansion */
```

The *f_basetype* member contains a null-terminated FSType name of the mounted target.

The following values can be returned in the *f_flag* field:

```
ST_RDONLY   0x01  /* read-only file system */
ST_NOSUID   0x02  /* does not support setuid/setgid semantics */
ST_NOTRUNC  0x04  /* does not truncate file names longer than
NAME_MAX */
```

The **fstatvfs()** function is similar to **statvfs()**, except that the file named by *path* in **statvfs()** is instead identified by an open file descriptor *filde*s obtained from a successful **open(2)**, **creat(2)**, **dup(2)**, **fcntl(2)**, or **pipe(2)** function call.

RETURN VALUES

Upon successful completion, 0 is returned. Otherwise, -1 is returned and **errno** is set to indicate the error.

ERRORS

The **statvfs()** and **fstatvfs()** functions will fail if:

EOverflow One of the values to be returned cannot be represented correctly in the structure pointed to by *buf*.

The **statvfs()** function will fail if:

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

EFAULT The *path* or *buf* argument points to an illegal address.

EINTR A signal was caught during the execution of the **statvfs()** function.

EIO An I/O error occurred while reading the file system.

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

ENAMETOOLONG The length of a *path* component exceeds **NAME_MAX** characters, or the length of *path* exceeds **PATH_MAX** characters.

ENOENT Either a component of the path prefix or the file referred to by *path* does not exist.

ENOLINK The *path* argument points to a remote machine and the link to that machine is no longer active.

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

The **fstatvfs()** function will fail if:

EBADF The *filde*s argument is not an open file descriptor.

EFAULT The *buf* argument points to an illegal address.

EINTR A signal was caught during the execution of the **fstatvfs()** function.

EIO An I/O error occurred while reading the file system.

USAGE The `statvfs()` and `fstatvfs()` functions have transitional interfaces for 64-bit file offsets. See `1f64(5)` .

SEE ALSO `chmod(2)` , `chown(2)` , `creat(2)` , `dup(2)` , `fcntl(2)` , `link(2)` , `mknod(2)` , `open(2)` , `pipe(2)` , `read(2)` , `time(2)` , `unlink(2)` , `utime(2)` , `write(2)` , `1f64(5)`

BUGS The values returned for `f_files` , `f_ffree` , and `f_favail` may not be valid for NFS mounted file systems.

NAME	stime – set system time and date
SYNOPSIS	<pre>#include <unistd.h> int stime(const time_t *tp);</pre>
DESCRIPTION	The stime() function sets the system's idea of the time and date. The <i>tp</i> argument points to the value of time as measured in seconds from 00:00:00 UTC January 1, 1970.
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	The stime() function will fail if: EINVAL The <i>tp</i> argument points to an invalid (negative) time value. EPERM The effective user of the calling process is not super-user.
SEE ALSO	<code>time(2)</code>

NAME	swapctl – manage swap space
SYNOPSIS	<pre>#include <sys/stat.h> #include <sys/swap.h> int swapctl(int cmd, void *arg);</pre>
DESCRIPTION	<p>The swapctl() function adds, deletes, or returns information about swap resources. <i>cmd</i> specifies one of the following options contained in <code><sys/swap.h></code>:</p> <pre>SC_ADD /* add a resource for swapping */ SC_LIST /* list the resources for swapping */ SC_REMOVE /* remove a resource for swapping */ SC_GETNSWP /* return number of swap resources */</pre> <p>When <code>SC_ADD</code> or <code>SC_REMOVE</code> is specified, <i>arg</i> is a pointer to a <code>swapres</code> structure containing the following members:</p> <pre>char *sr_name; /* pathname of resource */ off_t sr_start; /* offset to start of swap area */ off_t sr_length; /* length of swap area */</pre> <p>The <code>sr_start</code> and <code>sr_length</code> members are specified in 512-byte blocks. A swap resource can only be removed by specifying the same values for the <code>sr_start</code> and <code>sr_length</code> members as were specified when it was added. Swap resources need not be removed in the order in which they were added.</p> <p>When <code>SC_LIST</code> is specified, <i>arg</i> is a pointer to a <code>swaptable</code> structure containing the following members:</p> <pre>int swt_n; /* number of swapents following */ struct swapent swt_ent[]; /* array of swt_n swapents */</pre> <p>A <code>swapent</code> structure contains the following members:</p> <pre>char *ste_path; /* name of the swap file */ off_t ste_start; /* starting block for swapping */ off_t ste_length; /* length of swap area */ long ste_pages; /* number of pages for swapping */ long ste_free; /* number of ste_pages free */ long ste_flags; /* ST_INDEL bit set if swap file */ /* is now being deleted */</pre>

The `SC_LIST` function causes **swapctl()** to return at most `swt_n` entries. The return value of **swapctl()** is the number actually returned. The `ST_INDEL` bit is turned on in `ste_flags` if the swap file is in the process of being deleted.

When `SC_GETNSWP` is specified, **swapctl()** returns as its value the number of swap resources in use. `arg` is ignored for this operation.

The `SC_ADD` and `SC_REMOVE` functions will fail if calling process does not have appropriate privileges.

RETURN VALUES

Upon successful completion, the function **swapctl()** returns a value of 0 for `SC_ADD` or `SC_REMOVE`, the number of `struct swapent` entries actually returned for `SC_LIST`, or the number of swap resources in use for `SC_GETNSWP`. Upon failure, the function **swapctl()** returns a value of -1 and sets `errno` to indicate an error.

ERRORS

Under the following conditions, the function **swapctl()** fails and sets `errno` to:

EEXIST	Part of the range specified by <code>sr_start</code> and <code>sr_length</code> is already being used for swapping on the specified resource (<code>SC_ADD</code>).
EFAULT	Either <code>arg</code> , <code>sr_name</code> , or <code>ste_path</code> points to an illegal address.
EINVAL	The specified function value is not valid, the path specified is not a swap resource (<code>SC_REMOVE</code>), part of the range specified by <code>sr_start</code> and <code>sr_length</code> lies outside the resource specified (<code>SC_ADD</code>), or the specified swap area is less than one page (<code>SC_ADD</code>).
EISDIR	The path specified for <code>SC_ADD</code> is a directory.
ELOOP	Too many symbolic links were encountered in translating the pathname provided to <code>SC_ADD</code> or <code>SC_REMOVE</code> .
ENAMETOOLONG	The length of a component of the path specified for <code>SC_ADD</code> or <code>SC_REMOVE</code> exceeds <code>NAME_MAX</code> characters or the length of the path exceeds <code>PATH_MAX</code> characters and <code>_POSIX_NO_TRUNC</code> is in effect.
ENOENT	The pathname specified for <code>SC_ADD</code> or <code>SC_REMOVE</code> does not exist.

ENOMEM	An insufficient number of <code>struct swaptent</code> structures were provided to <code>SC_LIST</code> , or there were insufficient system storage resources available during an <code>SC_ADD</code> or <code>SC_REMOVE</code> , or the system would not have enough swap space after an <code>SC_REMOVE</code> .
ENOSYS	The pathname specified for <code>SC_ADD</code> or <code>SC_REMOVE</code> is not a file or block special device.
ENOTDIR	Pathname provided to <code>SC_ADD</code> or <code>SC_REMOVE</code> contained a component in the path prefix that was not a directory.
EPERM	The effective user of the calling process is not super-user.
EROFS	The pathname specified for <code>SC_ADD</code> is a read-only file system.
Additionally, the <code>swapctl()</code> function will fail for 32-bit interfaces if:	
EOVERFLOW	The amount of swap space configured on the machine is too large to be represented by a 32-bit quantity.

EXAMPLES

EXAMPLE 1 The usage of the `SC_GETNSWP` and `SC_LIST` commands.

The following example demonstrates the usage of the `SC_GETNSWP` and `SC_LIST` commands.

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

#define MAXSTRSIZE 80

main(argc, argv)
    int     argc;
    char    *argv[];
{
    swaptbl_t *s;
    int     i, n, num;
    char    *strtab; /* string table for path names */

again:
    if ((num = swapctl(SC_GETNSWP, 0)) == -1) {
        perror("swapctl: GETNSWP");
        exit(1);
    }
    if (num == 0) {
        fprintf(stderr, "No Swap Devices Configured\n");
    }
}
```

```

        exit(2);
    }
    /* allocate swaptable for num+1 entries */
    if ((s = (swaptbl_t *)
        malloc(num * sizeof(swapent_t) +
              sizeof(struct swaptable))) ==
        (void *) 0) {
        fprintf(stderr, "Malloc Failed\n");
        exit(3);
    }
    /* allocate num+1 string holders */
    if ((strtab = (char *)
        malloc((num + 1) * MAXSTRSIZE)) == (void *) 0) {
        fprintf(stderr, "Malloc Failed\n");
        exit(3);
    }
    /* initialize string pointers */
    for (i = 0; i < (num + 1); i++) {
        s->swt_ent[i].ste_path = strtab + (i * MAXSTRSIZE);
    }

    s->swt_n = num + 1;
    if ((n = swapctl(SC_LIST, s)) < 0) {
        perror("swapctl");
        exit(1);
    }
    if (n > num) {          /* more were added */
        free(s);
        free(strtab);
        goto again;
    }
    for (i = 0; i < n; i++)
        printf("%s %ld\n",
            s->swt_ent[i].ste_path, s->swt_ent[i].ste_pages);
}

```

NAME	symlink – make a symbolic link to a file														
SYNOPSIS	<pre>#include <unistd.h> int symlink(const char *<i>name1</i>, const char *<i>name2</i>);</pre>														
DESCRIPTION	<p>The symlink() function creates a symbolic link <i>name2</i> to the file <i>name1</i>. Either name may be an arbitrary pathname, the files need not be on the same file system, and <i>name1</i> may be nonexistent.</p> <p>The file to which the symbolic link points is used when an open(2) operation is performed on the link. A stat() operation performed on a symbolic link returns the linked-to file, while an lstat() operation returns information about the link itself. See stat(2). Unexpected results may occur when a symbolic link is made to a directory. To avoid confusion in applications, the readlink(2) call can be used to read the contents of a symbolic link.</p>														
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, errno is set to indicate the error, and the symbolic link is not made.														
ERRORS	<table border="0"> <tr> <td style="vertical-align: top;">EACCES</td> <td>Search permission is denied for a component of the path prefix of <i>name2</i>.</td> </tr> <tr> <td style="vertical-align: top;">EDQUOT</td> <td>The directory where the entry for the new symbolic link is being placed cannot be extended because the user's quota of disk blocks on that file system has been exhausted; the new symbolic link cannot be created because the user's quota of disk blocks on that file system has been exhausted; or the user's quota of inodes on the file system where the file is being created has been exhausted.</td> </tr> <tr> <td style="vertical-align: top;">EEXIST</td> <td>The file referred to by <i>name2</i> already exists.</td> </tr> <tr> <td style="vertical-align: top;">EFAULT</td> <td>The <i>name1</i> or <i>name2</i> argument points to an illegal address.</td> </tr> <tr> <td style="vertical-align: top;">EIO</td> <td>An I/O error occurs while reading from or writing to the file system.</td> </tr> <tr> <td style="vertical-align: top;">ELOOP</td> <td>Too many symbolic links are encountered in translating <i>name2</i>.</td> </tr> <tr> <td style="vertical-align: top;">ENAMETOOLONG</td> <td>The length of the <i>name2</i> argument exceeds PATH_MAX, or the length of a <i>name2</i> component</td> </tr> </table>	EACCES	Search permission is denied for a component of the path prefix of <i>name2</i> .	EDQUOT	The directory where the entry for the new symbolic link is being placed cannot be extended because the user's quota of disk blocks on that file system has been exhausted; the new symbolic link cannot be created because the user's quota of disk blocks on that file system has been exhausted; or the user's quota of inodes on the file system where the file is being created has been exhausted.	EEXIST	The file referred to by <i>name2</i> already exists.	EFAULT	The <i>name1</i> or <i>name2</i> argument points to an illegal address.	EIO	An I/O error occurs while reading from or writing to the file system.	ELOOP	Too many symbolic links are encountered in translating <i>name2</i> .	ENAMETOOLONG	The length of the <i>name2</i> argument exceeds PATH_MAX , or the length of a <i>name2</i> component
EACCES	Search permission is denied for a component of the path prefix of <i>name2</i> .														
EDQUOT	The directory where the entry for the new symbolic link is being placed cannot be extended because the user's quota of disk blocks on that file system has been exhausted; the new symbolic link cannot be created because the user's quota of disk blocks on that file system has been exhausted; or the user's quota of inodes on the file system where the file is being created has been exhausted.														
EEXIST	The file referred to by <i>name2</i> already exists.														
EFAULT	The <i>name1</i> or <i>name2</i> argument points to an illegal address.														
EIO	An I/O error occurs while reading from or writing to the file system.														
ELOOP	Too many symbolic links are encountered in translating <i>name2</i> .														
ENAMETOOLONG	The length of the <i>name2</i> argument exceeds PATH_MAX , or the length of a <i>name2</i> component														

	exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.
ENOENT	A component of the path prefix of <i>name2</i> does not exist.
ENOSPC	The directory in which the entry for the new symbolic link is being placed cannot be extended because no space is left on the file system containing the directory; the new symbolic link cannot be created because no space is left on the file system which will contain the link; or there are no free inodes on the file system on which the file is being created.
ENOSYS	The file system does not support symbolic links
ENOTDIR	A component of the path prefix of <i>name2</i> is not a directory.
EROFS	The file <i>name2</i> would reside on a read-only file system.

SEE ALSO `cp(1)`, `link(2)`, `open(2)`, `readlink(2)`, `stat(2)`, `unlink(2)`

NAME	sync – update super block
SYNOPSIS	<pre>#include <unistd.h> void sync(void);</pre>
DESCRIPTION	<p>The sync() function writes all information in memory that should be on disk, including modified super blocks, modified inodes, and delayed block I/O.</p> <p>Unlike fsync(3C), which completes the writing before it returns, sync() schedules but does not necessarily complete the writing before returning.</p>
USAGE	The sync() function should be used by applications that examine a file system, such as fsck(1M) , and df(1M) , and is mandatory before rebooting.
SEE ALSO	df(1M) , fsck(1M) , fsync(3C)

NAME	sysfs – get file system type information
SYNOPSIS	<pre>#include <sys/fstyp.h> #include <sys/fsid.h> int sysfs(int opcode, const char *fsname); int sysfs(int opcode, int fs_index, char *buf); int sysfs(int opcode);</pre>
DESCRIPTION	<p>The sysfs() function returns information about the file system types configured in the system. The number of arguments accepted by sysfs() depends on the <i>opcode</i> argument, which can take the following values:</p> <p>GETFSIND Translate <i>fsname</i>, a null-terminated file-system type identifier, into a file-system type index.</p> <p>GETFSTYP Translate <i>fs_index</i>, a file-system type index, into a null-terminated file-system type identifier and write it into the buffer pointed to by <i>buf</i>, which must be at least of size <code>FSTYPSZ</code> as defined in <code><sys/fstyp.h></code>.</p> <p>GETNFSTYP Return the total number of file system types configured in the system.</p>
RETURN VALUES	<p>Upon successful completion, the value returned depends upon the <i>opcode</i> argument as follows:</p> <p>GETFSIND the file-system type index</p> <p>GETFSTYP 0</p> <p>GETNFSTYP the number of file system types configured</p> <p>Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.</p>
ERRORS	<p>The sysfs() function will fail if:</p> <p>EFAULT The <i>buf</i> or <i>fsname</i> argument points to an illegal address.</p> <p>EINVAL The <i>fsname</i> argument points to an invalid file-system identifier; the <i>fs_index</i> argument is 0 or invalid; or the <i>opcode</i> argument is invalid.</p>

NAME	sysinfo – get and set system information strings								
SYNOPSIS	<pre>#include <sys/systeminfo.h> long sysinfo(int <i>command</i>, char *<i>buf</i>, long <i>count</i>);</pre>								
DESCRIPTION	<p>The sysinfo() function copies information relating to the operating system on which the process is executing into the buffer pointed to by <i>buf</i>. It can also set certain information where appropriate <i>commands</i> are available. <i>count</i> is the size of the buffer.</p> <p>The POSIX P1003.1 interface (see standards(5)) sysconf(3C) provides a similar class of configuration information, but returns an integer rather than a string.</p> <p>The <i>commands</i> available are:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">SI_SYSNAME</td> <td>Copy into the array pointed to by <i>buf</i> the string that would be returned by uname(2) in the <i>sysname</i> field. This is the name of the implementation of the operating system, for example, SunOS or UTS.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">SI_HOSTNAME</td> <td>Copy into the array pointed to by <i>buf</i> a string that names the present host machine. This is the string that would be returned by uname(2) in the <i>nodename</i> field. This hostname or nodename is often the name the machine is known by locally. The <i>hostname</i> is the name of this machine as a node in some network. Different networks may have different names for the node, but presenting the nodename to the appropriate network directory or name-to-address mapping service should produce a transport end point address. The name may not be fully qualified. Internet host names may be up to 256 bytes in length (plus the terminating null).</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">SI_SET_HOSTNAME</td> <td>Copy the null-terminated contents of the array pointed to by <i>buf</i> into the string maintained by the kernel whose value will be returned by succeeding calls to sysinfo() with the command SI_HOSTNAME. This command requires that the effective-user-id be super-user.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">SI_RELEASE</td> <td>Copy into the array pointed to by <i>buf</i> the string that would be returned by uname(2) in the <i>release</i> field. Typical values might be 5.2 or 4.1.</td> </tr> </table>	SI_SYSNAME	Copy into the array pointed to by <i>buf</i> the string that would be returned by uname(2) in the <i>sysname</i> field. This is the name of the implementation of the operating system, for example, SunOS or UTS.	SI_HOSTNAME	Copy into the array pointed to by <i>buf</i> a string that names the present host machine. This is the string that would be returned by uname(2) in the <i>nodename</i> field. This hostname or nodename is often the name the machine is known by locally. The <i>hostname</i> is the name of this machine as a node in some network. Different networks may have different names for the node, but presenting the nodename to the appropriate network directory or name-to-address mapping service should produce a transport end point address. The name may not be fully qualified. Internet host names may be up to 256 bytes in length (plus the terminating null).	SI_SET_HOSTNAME	Copy the null-terminated contents of the array pointed to by <i>buf</i> into the string maintained by the kernel whose value will be returned by succeeding calls to sysinfo() with the command SI_HOSTNAME. This command requires that the effective-user-id be super-user.	SI_RELEASE	Copy into the array pointed to by <i>buf</i> the string that would be returned by uname(2) in the <i>release</i> field. Typical values might be 5.2 or 4.1.
SI_SYSNAME	Copy into the array pointed to by <i>buf</i> the string that would be returned by uname(2) in the <i>sysname</i> field. This is the name of the implementation of the operating system, for example, SunOS or UTS.								
SI_HOSTNAME	Copy into the array pointed to by <i>buf</i> a string that names the present host machine. This is the string that would be returned by uname(2) in the <i>nodename</i> field. This hostname or nodename is often the name the machine is known by locally. The <i>hostname</i> is the name of this machine as a node in some network. Different networks may have different names for the node, but presenting the nodename to the appropriate network directory or name-to-address mapping service should produce a transport end point address. The name may not be fully qualified. Internet host names may be up to 256 bytes in length (plus the terminating null).								
SI_SET_HOSTNAME	Copy the null-terminated contents of the array pointed to by <i>buf</i> into the string maintained by the kernel whose value will be returned by succeeding calls to sysinfo() with the command SI_HOSTNAME. This command requires that the effective-user-id be super-user.								
SI_RELEASE	Copy into the array pointed to by <i>buf</i> the string that would be returned by uname(2) in the <i>release</i> field. Typical values might be 5.2 or 4.1.								

SI_VERSION	Copy into the array pointed to by <i>buf</i> the string that would be returned by <code>uname(2)</code> in the <i>version</i> field. The syntax and semantics of this string are defined by the system provider.
SI_MACHINE	Copy into the array pointed to by <i>buf</i> the string that would be returned by <code>uname(2)</code> in the <i>machine</i> field, for example, <i>sun4c</i> , <i>sun4d</i> , or <i>sun4m</i> .
SI_ARCHITECTURE	Copy into the array pointed to by <i>buf</i> a string describing the basic instruction set architecture of the current system, for example, <i>sparc</i> , <i>m68030</i> , <i>m32100</i> , or <i>i386</i> . These names may not match predefined names in the C language compilation system.
SI_ISALIST	<p>Copy into the array pointed to by <i>buf</i> the names of the variant instruction set architectures executable on the current system.</p> <p>The names are space-separated and are ordered in the sense of best performance. That is, earlier-named instruction sets may contain more instructions than later-named instruction sets; a program that is compiled for an earlier-named instruction set will most likely run faster on this machine than the same program compiled for a later-named instruction set.</p> <p>Programs compiled for an instruction set that does not appear in the list will most likely experience performance degradation or not run at all on this machine.</p> <p>The instruction set names known to the system are listed in <code>isalist(5)</code>; these names may or may not match predefined names or compiler options in the C language compilation system.</p>
SI_PLATFORM	Copy into the array pointed to by <i>buf</i> a string describing the specific model of the hardware platform, for example, <i>SUNW,Sun_4_75</i> , <i>SUNW,SPARCsystem-600</i> , or <i>i86pc</i> .

SI_HW_PROVIDER	Copies the name of the hardware manufacturer into the array pointed to by <i>buf</i> .
SI_HW_SERIAL	Copy into the array pointed to by <i>buf</i> a string which is the ASCII representation of the hardware-specific serial number of the physical machine on which the function is executed. Note that this may be implemented in Read-Only Memory, using software constants set when building the operating system, or by other means, and may contain non-numeric characters. It is anticipated that manufacturers will not issue the same “serial number” to more than one physical machine. The pair of strings returned by SI_HW_PROVIDER and SI_HW_SERIAL is likely to be unique across all vendor’s SVR4 implementations.
SI_SRPC_DOMAIN	Copies the Secure Remote Procedure Call domain name into the array pointed to by <i>buf</i> .
SI_SET_SRPC_DOMAIN	Set the string to be returned by sysinfo() with the SI_SRPC_DOMAIN command to the value contained in the array pointed to by <i>buf</i> . This command requires that the effective-user-id be super-user.

RETURN VALUES

Upon successful completion, the value returned indicates the buffer size in bytes required to hold the complete value and the terminating null character. If this value is no greater than the value passed in *count*, the entire string was copied. If this value is greater than *count*, the string copied into *buf* has been truncated to *count* - 1 bytes plus a terminating null character.

Otherwise, -1 is returned and *errno* is set to indicate the error.

ERRORS

The **sysinfo()** function will fail if:

EFAULT	The <i>buf</i> argument does not point to a valid address.
EINVAL	The data for a SET command exceeds the limits established by the implementation.
EPERM	The effective user of the calling process is not super-user.

USAGE

In many cases there is no corresponding programmatic interface to set these values; such strings are typically settable only by the system administrator modifying entries in the `/etc/system` directory or the code provided by the particular OEM reading a serial number or code out of read-only memory, or hard-coded in the version of the operating system.

A good starting guess for *count* is 257, which is likely to cover all strings returned by this interface in typical installations.

SEE ALSO

`uname(2)`, `gethostid(3C)`, `gethostname(3C)`, `sysconf(3C)`, `isalist(5)`, `standards(5)`

NAME time – get time

SYNOPSIS

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

time_t time(time_t *tloc);
```

DESCRIPTION The **time()** function returns the value of time in seconds since 00:00:00 UTC, January 1, 1970.

If *tloc* is non-zero, the return value is also stored in the location to which *tloc* points. If *tloc* points to an illegal address, **time()** fails and its actions are undefined.

RETURN VALUES Upon successful completion, **time()** returns the value of time. Otherwise, `(time_t)-1` is returned and `errno` is set to indicate the error.

ATTRIBUTES See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO **stime(2)**, **ctime(3C)**, **attributes(5)**

NAME	times – get process and child process times
SYNOPSIS	<pre>#include <sys/times.h> #include <limits.h> clock_t times(struct tms *buffer);</pre>
DESCRIPTION	<p>The times() function fills the <code>tms</code> structure pointed to by <i>buffer</i> with time-accounting information. The <code>tms</code> structure, defined in <code><sys/times.h></code>, contains the following members:</p> <pre>clock_t tms_utime; clock_t tms_stime; clock_t tms_cutime; clock_t tms_cstime;</pre> <p>All times are reported in clock ticks. The specific value for a clock tick is defined by the variable <code>CLK_TCK</code>, found in the header <code><limits.h></code>.</p> <p>The times of a terminated child process are included in the <code>tms_cutime</code> and <code>tms_cstime</code> members of the parent when <code>wait(2)</code> or <code>waitpid(2)</code> returns the process ID of this terminated child. If a child process has not waited for its children, their times will not be included in its times.</p> <p>The <code>tms_utime</code> member is the CPU time used while executing instructions in the user space of the calling process.</p> <p>The <code>tms_stime</code> member is the CPU time used by the system on behalf of the calling process.</p> <p>The <code>tms_cutime</code> member is the sum of the <code>tms_utime</code> and the <code>tms_cutime</code> of the child processes.</p> <p>The <code>tms_cstime</code> member is the sum of the <code>tms_stime</code> and the <code>tms_cstime</code> of the child processes.</p>
RETURN VALUES	<p>Upon successful completion, times() returns the elapsed real time, in clock ticks, since an arbitrary point in the past (for example, system start-up time). This point does not change from one invocation of times() within the process to another. The return value may overflow the possible range of type <code>clock_t</code>. If times() fails, <code>(clock_t)-1</code> is returned and <code>errno</code> is set to indicate the error.</p>
ERRORS	<p>The times() function will fail if:</p> <p>EFAULT The <i>buffer</i> argument points to an illegal address.</p>
ATTRIBUTES	<p>See <code>attributes(5)</code> for descriptions of the following attributes:</p>

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

`time(1)`, `timex(1)`, `exec(2)`, `fork(2)`, `time(2)`, `wait(2)`, `waitid(2)`,
`waitpid(2)`, `attributes(5)`

NAME	uadmin – administrative control																				
SYNOPSIS	<pre>#include <sys/uadmin.h> int uadmin(int cmd, int fcn, int mdep);</pre>																				
DESCRIPTION	<p>The uadmin() function provides control for basic administrative functions. This function is tightly coupled to the system administrative procedures and is not intended for general use. The argument <i>mdep</i> is provided for machine-dependent use and is not defined here.</p> <p>As specified by <i>cmd</i>, the following commands are available:</p> <table border="0"> <tr> <td style="vertical-align: top;">A_SHUTDOWN</td> <td>The system is shut down. All user processes are killed, the buffer cache is flushed, and the root file system is unmounted. The action to be taken after the system has been shut down is specified by <i>fcn</i>. The functions are generic; the hardware capabilities vary on specific machines.</td> </tr> <tr> <td style="vertical-align: top;">AD_HALT</td> <td>Halt the processor(s).</td> </tr> <tr> <td style="vertical-align: top;">AD_POWEROFF</td> <td>Halt the processor(s) and turn off the power.</td> </tr> <tr> <td style="vertical-align: top;">AD_BOOT</td> <td>Reboot the system, using the kernel file.</td> </tr> <tr> <td style="vertical-align: top;">AD_IBOOT</td> <td>Interactive reboot; user is prompted for bootable program name.</td> </tr> <tr> <td style="vertical-align: top;">A_REBOOT</td> <td>The system stops immediately without any further processing. The action to be taken next is specified by <i>fcn</i> as above.</td> </tr> <tr> <td style="vertical-align: top;">A_REMOUNT</td> <td>The root file system is mounted again after having been fixed. This should be used only during the startup process.</td> </tr> <tr> <td style="vertical-align: top;">A_FREEZE</td> <td>Suspend the whole system. The system state is preserved in the state file. The following three subcommands are available.</td> </tr> <tr> <td style="vertical-align: top;">AD_COMPRESS</td> <td>Save the system state to the state file with compression of data.</td> </tr> <tr> <td style="vertical-align: top;">AD_CHECK</td> <td>Check if your system supports suspend and resume. Without performing a system suspend/resume, this command checks if this feature is currently available on your system.</td> </tr> </table>	A_SHUTDOWN	The system is shut down. All user processes are killed, the buffer cache is flushed, and the root file system is unmounted. The action to be taken after the system has been shut down is specified by <i>fcn</i> . The functions are generic; the hardware capabilities vary on specific machines.	AD_HALT	Halt the processor(s).	AD_POWEROFF	Halt the processor(s) and turn off the power.	AD_BOOT	Reboot the system, using the kernel file.	AD_IBOOT	Interactive reboot; user is prompted for bootable program name.	A_REBOOT	The system stops immediately without any further processing. The action to be taken next is specified by <i>fcn</i> as above.	A_REMOUNT	The root file system is mounted again after having been fixed. This should be used only during the startup process.	A_FREEZE	Suspend the whole system. The system state is preserved in the state file. The following three subcommands are available.	AD_COMPRESS	Save the system state to the state file with compression of data.	AD_CHECK	Check if your system supports suspend and resume. Without performing a system suspend/resume, this command checks if this feature is currently available on your system.
A_SHUTDOWN	The system is shut down. All user processes are killed, the buffer cache is flushed, and the root file system is unmounted. The action to be taken after the system has been shut down is specified by <i>fcn</i> . The functions are generic; the hardware capabilities vary on specific machines.																				
AD_HALT	Halt the processor(s).																				
AD_POWEROFF	Halt the processor(s) and turn off the power.																				
AD_BOOT	Reboot the system, using the kernel file.																				
AD_IBOOT	Interactive reboot; user is prompted for bootable program name.																				
A_REBOOT	The system stops immediately without any further processing. The action to be taken next is specified by <i>fcn</i> as above.																				
A_REMOUNT	The root file system is mounted again after having been fixed. This should be used only during the startup process.																				
A_FREEZE	Suspend the whole system. The system state is preserved in the state file. The following three subcommands are available.																				
AD_COMPRESS	Save the system state to the state file with compression of data.																				
AD_CHECK	Check if your system supports suspend and resume. Without performing a system suspend/resume, this command checks if this feature is currently available on your system.																				

AD_FORCE Force **AD_COMPRESS** even when threads of drivers are not suspendable.

RETURN VALUES

Upon successful completion, the value returned depends on *cmd* as follows:

A_SHUTDOWN Never returns.

A_REBOOT Never returns.

A_FREEZE 0 upon resume.

A_REMOUNT 0.

Otherwise, -1 is returned and **errno** is set to indicate the error.

ERRORS

The **uadmin()** function will fail if:

EPERM The effective user of the calling process is not super-user.

ENOMEM Suspend/resume ran out of physical memory.

ENOSPC Suspend/resume could not allocate enough space on the root file system to store system information.

ENOTSUP Suspend/resume not supported on this platform.

ENXIO Unable to successfully suspend system.

EBUSY Suspend already in progress.

SEE ALSO

kernel(1M), **uadmin(1M)**

NAME	ulimit – get and set process limits
SYNOPSIS	<pre>#include <ulimit.h> long ulimit(int cmd, /* newlimit */...);</pre>
DESCRIPTION	<p>The ulimit() function provides for control over process limits. It is effective in limiting the growth of regular files. Pipes are limited to <code>PIPE_MAX</code> bytes.</p> <p>The <i>cmd</i> values, defined in <code><ulimit.h></code>, include:</p> <p>UL_GETFSIZE Return the soft file size limit of the process. The limit is in units of 512-byte blocks and is inherited by child processes. Files of any size can be read. The return value is the integer part of the soft file size limit divided by 512. If the result cannot be represented as a <code>long int</code>, the result is unspecified.</p> <p>UL_SETFSIZE Set the hard and soft file size limits for output operations of the process to the value of the second argument, taken as a <code>long int</code>. Any process may decrease its own hard limit, but only a process with appropriate privileges may increase the limit. The new file size limit is returned. The hard and soft file size limits are set to the specified value multiplied by 512. If the result would overflow an <code>rlimit_t</code>, the actual value set is unspecified.</p> <p>UL_GMEMLIM Get the maximum possible break value (see <code>brk(2)</code>).</p> <p>UL_GDESLIM Get the current value of the maximum number of open files per process configured in the system.</p>
RETURN VALUES	Upon successful completion, ulimit() returns the value of the requested limit. Otherwise, <code>-1</code> is returned, the limit is not changed, and <code>errno</code> is set to indicate the error.
ERRORS	<p>The ulimit() function will fail if:</p> <p>EINVAL The <i>cmd</i> argument is not valid.</p> <p>EPERM A process not having appropriate privileges attempts to increase its file size limit.</p>
USAGE	Since all return values are permissible in a successful situation, an application wishing to check for error situations should set <code>errno</code> to 0, then call ulimit() , and if it returns <code>-1</code> , check if <code>errno</code> is non-zero.

The **getrlimit()** and **setrlimit()** functions provide a more general interface for controlling process limits, and are preferred over **ulimit()**. See **getrlimit(2)**.

SEE ALSO

brk(2), **getrlimit(2)**, **write(2)**

NAME | umask – set and get file creation mask

SYNOPSIS | `#include <sys/types.h>`
| `#include <sys/stat.h>`
| `mode_t umask(mode_t cmask);`

DESCRIPTION | The **umask()** function sets the process's file mode creation mask to *cmask* and returns the previous value of the mask. Only the access permission bits of *cmask* and the file mode creation mask are used. The mask is inherited by child processes. See **intro(2)** for more information on masks.

RETURN VALUES | The previous value of the file mode creation mask is returned.

ATTRIBUTES | See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO | **mkdir(1)**, **sh(1)**, **intro(2)**, **chmod(2)**, **creat(2)**, **mknod(2)**, **open(2)**, **attributes(5)**, **stat(5)**

NAME	umount – unmount a file system
SYNOPSIS	<pre>#include <sys/mount.h> int umount(const char *file);</pre>
DESCRIPTION	The umount() function requests that a previously mounted file system contained on the block special device or directory identified by <i>file</i> be unmounted. The <i>file</i> argument is a pointer to a path name. After unmounting the file system, the directory upon which the file system was mounted reverts to its ordinary interpretation.
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The umount() function will fail if:</p> <p>EBUSY A file on <i>file</i> is busy.</p> <p>EFAULT The file pointed to by <i>file</i> points to an illegal address.</p> <p>EINVAL The file pointed to by <i>file</i> is not mounted.</p> <p>ENOENT The file pointed to by <i>file</i> does not exist.</p> <p>ELOOP Too many symbolic links were encountered in translating the path pointed to by <i>file</i>.</p> <p>ENAMETOOLONG The length of the <i>file</i> argument exceeds <code>PATH_MAX</code>, or the length of a <i>file</i> component exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.</p> <p>ENOLINK The file pointed to by <i>file</i> is on a remote machine, and the link to that machine is no longer active.</p> <p>ENOTBLK The file pointed to by <i>file</i> is not a block special device.</p> <p>EPERM The process's effective user ID is not super-user.</p> <p>EREMOTE The file pointed to by <i>file</i> is remote.</p>
USAGE	The umount() function may be invoked only by the super-user.
SEE ALSO	<code>mount(2)</code>

NAME	uname – get name of current operating system				
SYNOPSIS	<pre>#include <sys/utsname.h> int uname(struct utsname *name);</pre>				
DESCRIPTION	<p>The uname() function stores information identifying the current operating system in the structure pointed to by <i>name</i>.</p> <p>The uname() function uses the <code>utsname</code> structure, defined in <code><sys/utsname.h></code>, whose members include:</p> <pre>char sysname[SYS_NMLN]; char odename[SYS_NMLN]; char release[SYS_NMLN]; char version[SYS_NMLN]; char machine[SYS_NMLN];</pre> <p>The uname() function returns a null-terminated character string naming the current operating system in the character array <code>sysname</code>. Similarly, the <code>nodename</code> member contains the name by which the system is known on a communications network. The <code>release</code> and <code>version</code> members further identify the operating system. The <code>machine</code> member contains a standard name that identifies the hardware on which the operating system is running.</p>				
RETURN VALUES	Upon successful completion, a non-negative value is returned. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.				
ERRORS	<p>The uname() function will fail if:</p> <p>EFAULT The <i>name</i> argument points to an illegal address.</p>				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Async-Signal-Safe				
SEE ALSO	uname(1) , sysinfo(2) , sysconf(3C) , attributes(5)				

NAME	unlink – remove directory entry
SYNOPSIS	<pre>#include <unistd.h> int unlink(const char *path);</pre>
DESCRIPTION	<p>The unlink() function removes a link to a file. If <i>path</i> names a symbolic link, unlink() removes the symbolic link named by <i>path</i> and does not affect any file or directory named by the contents of the symbolic link. Otherwise, unlink() removes the link named by the pathname pointed to by <i>path</i> and decrements the link count of the file referenced by the link.</p> <p>When the file's link count becomes 0 and no process has the file open, the space occupied by the file will be freed and the file will no longer be accessible. If one or more processes have the file open when the last link is removed, the link will be removed before unlink() returns, but the removal of the file contents will be postponed until all references to the file are closed.</p> <p>The <i>path</i> argument must not name a directory unless the process has appropriate privileges and the implementation supports using unlink() on directories.</p> <p>Upon successful completion, unlink() will mark for update the <code>st_ctime</code> and <code>st_mtime</code> fields of the parent directory. If the file's link count is not 0, the <code>st_ctime</code> field of the file will be marked for update.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, <code>errno</code> is set to indicate the error, and the file is not unlinked.
ERRORS	<p>The unlink() function will fail if:</p> <p>EACCES Search permission is denied for a component of the <i>path</i> prefix; write permission is denied on the directory containing the link to be removed; the parent directory has the sticky bit set and the file is not writable by the user; or the user does not own the parent directory and the user does not own the file.</p> <p>EBUSY The entry to be unlinked is the mount point for a mounted file system.</p> <p>EFAULT The <i>path</i> argument points to an illegal address.</p> <p>EINTR A signal was caught during the execution of the unlink() function.</p> <p>ELOOP Too many symbolic links were encountered in translating <i>path</i>.</p>

ENAMETOOLONG The length of the *path* argument exceeds `PATH_MAX`, or the length of a *path* component exceeds `NAME_MAX` while `_POSIX_NO_TRUNC` is in effect.

ENOENT The named file does not exist or is a null pathname.

ENOLINK The *path* argument points to a remote machine and the link to that machine is no longer active.

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

EPERM The named file is a directory and the effective user of the calling process is not super-user.

EROFS The directory entry to be unlinked is part of a read-only file system.

The `unlink()` function may fail if:

ENAMETOOLONG Pathname resolution of a symbolic link produced an intermediate result whose length exceeds `PATH_MAX`.

ETXTBSY The entry to be unlinked is the last directory entry to a pure procedure (shared text) file that is being executed.

USAGE Applications should use `rmdir(2)` to remove a directory.

ATTRIBUTES See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO `rm(1)`, `close(2)`, `link(2)`, `open(2)`, `rmdir(2)`, `remove(3C)`, `attributes(5)`

NAME	ustat – get file system statistics
SYNOPSIS	<pre>#include <sys/types.h> #include <ustat.h> int ustat(dev_t dev, struct ustat *buf);</pre>
DESCRIPTION	<p>The ustat() function returns information about a mounted file system. The <i>dev</i> argument is a device number identifying a device containing a mounted file system (see makedev(3C)). The <i>buf</i> argument is a pointer to a ustat structure that includes the following members:</p> <pre>daddr_t f_tfree; /* Total free blocks */ ino_t f_tinode; /* Number of free inodes */ char f_fname[6]; /* Filsys name */ char f_fpack[6]; /* Filsys pack name */</pre> <p>The <i>f_fname</i> and <i>f_fpack</i> members may not contain significant information on all systems; in this case, these members will contain the null character as the first character.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <i>errno</i> is set to indicate the error.
ERRORS	<p>The ustat() function will fail if:</p> <p>ECOMM The <i>dev</i> argument is on a remote machine and the link to that machine is no longer active.</p> <p>EFAULT The <i>buf</i> argument points to an illegal address.</p> <p>EINTR A signal was caught during the execution of the ustat() function.</p> <p>EINVAL The <i>dev</i> argument is not the device number of a device containing a mounted file system.</p> <p>ENOLINK The <i>dev</i> argument refers to a device on a remote machine and the link to that machine is no longer active.</p>
USAGE	The statvfs(2) function should be used in favor of ustat() .
SEE ALSO	stat(2) , statvfs(2) , makedev(3C)
BUGS	The NFS revision 2 protocol does not permit the number of free files to be provided to the client; therefore, when ustat() has completed on an NFS file system, <i>f_tinode</i> is always -1.

NAME	utime – set file access and modification times
SYNOPSIS	<pre>#include <sys/types.h> #include <utime.h> int utime(const char *path, const struct utimbuf *times);</pre>
DESCRIPTION	<p>The utime() function sets the access and modification times of the file pointed to by <i>path</i>, and causes the time of the last file status change (<i>st_ctime</i>) to be updated.</p> <p>If <i>times</i> is <code>NULL</code>, the access and modification times of the file are set to the current time. A process must be the owner of the file or have write permission to use utime() in this manner.</p> <p>If <i>times</i> is not <code>NULL</code>, <i>times</i> is interpreted as a pointer to a <code>utimbuf</code> structure (defined in <code><utime.h></code>) and the access and modification times are set to the values contained in the designated structure. Only the owner of the file or the super-user may use utime() in this manner.</p> <p>The <code>utimbuf</code> structure contains the following members:</p> <pre>time_t actime; /* access time */ time_t modtime; /* modification time */</pre> <p>The times contained in the members of the <code>utimbuf</code> structure are measured in seconds since 00:00:00 UTC, January 1, 1970.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The utime() function will fail if:</p> <p>EACCES Search permission is denied by a component of the <i>path</i> prefix; or 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 <i>times</i> is <code>NULL</code>.</p> <p>EFAULT The <i>path</i> argument points to an illegal address.</p> <p>EINTR A signal was caught during the execution of the utime() function.</p> <p>EIO An I/O error occurred while reading from or writing to the file system.</p>

ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
ENAMETOOLONG	The length of the <i>path</i> argument exceeds <code>PATH_MAX</code> , or the length of a <i>path</i> component exceeds <code>NAME_MAX</code> while <code>_POSIX_NO_TRUNC</code> is in effect.
ENOENT	The named file does not exist or is a null pathname.
ENOLINK	The <i>path</i> argument points to a remote machine and the link to that machine is no longer active.
ENOTDIR	A component of the <i>path</i> prefix is not a directory.
EPERM	The effective user of the calling process is not super-user and not the owner of the file, and <i>times</i> is not <code>NULL</code> .
EROFS	The file system containing the file is mounted read-only.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

`stat(2)`, `attributes(5)`

NAME	utimes – set file access and modification times
SYNOPSIS	<pre>#include <sys/time.h> int utimes(const char *path, const struct timeval times);</pre>
DESCRIPTION	<p>The utimes() function sets the access and modification times of the file pointed to by the <i>path</i> argument to the value of the <i>times</i> argument. It allows time specifications accurate to the microsecond.</p> <p>The <i>times</i> argument is an array of <code>timeval</code> structures. The first array member represents the date and time of last access, and the second member represents the date and time of last modification. The times in the <code>timeval</code> structure are measured in seconds and microseconds since the Epoch, although rounding toward the nearest second may occur.</p> <p>If the <i>times</i> argument is a null pointer, the access and modification times of the file are set to the current time. The effective user ID of the process must be the same as the owner of the file, or must have write access to the file or super-user privileges to use this call in this manner. Upon completion, utimes() will mark the time of the last file status change, <code>st_ctime</code>, for update.</p>
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, <code>errno</code> is set to indicate the error, and the file times will not be affected.
ERRORS	<p>The utimes() function will fail if:</p> <p>EACCES Search permission is denied by a component of the path prefix; or the <i>times</i> argument is a null pointer and the effective user ID of the process does not match the owner of the file and write access is denied.</p> <p>EFAULT The <i>path</i> or <i>times</i> argument points to an illegal address.</p> <p>EINTR A signal was caught during the execution of the utimes() function.</p> <p>EINVAL The number of microseconds specified in one or both of the <code>timeval</code> structures pointed to by <i>times</i> was greater than or equal to 1,000,000 or less than 0.</p> <p>EIO An I/O error occurred while reading from or writing to the file system.</p> <p>ELOOP Too many symbolic links were encountered in resolving <i>path</i>.</p>

ENAMETOOLONG	The length of the <i>path</i> argument exceeds <code>PATH_MAX</code> or a pathname component is longer than <code>NAME_MAX</code> .
ENOLINK	The <i>path</i> argument points to a remote machine and the link to that machine is no longer active.
ENOENT	A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.
ENOTDIR	A component of the path prefix is not a directory.
EPERM	The <i>times</i> argument is not a null pointer and the calling process's effective user ID has write access to the file but does not match the owner of the file and the calling process does not have the appropriate privileges.
EROFS	The file system containing the file is read-only.
The utimes() function may fail if:	
ENAMETOOLONG	Path name resolution of a symbolic link produced an intermediate result whose length exceeds <code>PATH_MAX</code> .

SEE ALSO**stat(2)**

NAME	vfork – spawn new process in a virtual memory efficient way
SYNOPSIS	<pre>#include <unistd.h> pid_t vfork(void);</pre>
DESCRIPTION	<p>The vfork() function creates new processes without fully copying the address space of the old process. This function is useful in instances where the purpose of a fork(2) operation would be to create a new system context for an execve() operation (see exec(2)).</p> <p>Unlike with the fork() function, the child process borrows the parent's memory and thread of control until a call to execve() or an exit (either abnormally or by a call to _exit() (see exit(2)). The parent process is suspended while the child is using its resources.</p> <p>In a multithreaded application, vfork() borrows only the thread of control that called vfork() in the parent; that is, the child contains only one thread. In that sense, vfork() behaves like fork().</p> <p>The vfork() function can normally be used the same way as fork(). The procedure that called vfork(), however, should not return while running in the child's context, since the eventual return from vfork() would be to a stack frame that no longer exists. The _exit() function should be used in favor of exit(3C) if unable to perform an execve() operation, since exit() will flush and close standard I/O channels, and thereby corrupt the parent process's standard I/O data structures. The _exit() function should be used even with fork() to avoid flushing the buffered data twice.</p>
RETURN VALUES	Upon successful completion, vfork() returns 0 to the child process and returns the process ID of the child process to the parent process. Otherwise, -1 is returned to the parent process, no child process is created, and errno is set to indicate the error.
ERRORS	<p>The vfork() function will fail if:</p> <p>EAGAIN The system-imposed limit on the total number of processes under execution (either system-quality or by a single user) would be exceeded. This limit is determined when the system is generated.</p> <p>ENOMEM There is insufficient swap space for the new process.</p>
SEE ALSO	exec(2) , exit(2) , fork(2) , ioctl(2) , wait(2) , exit(3C)
NOTES	The use of vfork() for any purpose other than as a prelude to an immediate call to a function from the exec family or to _exit() is not advised.

The **vfork()** function is unsafe in multithreaded applications.

This function will be eliminated in a future release. 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 `ioctl`s are allowed and input attempts result in an EOF indication.

On some systems, the implementation of **vfork()** causes the parent to inherit register values from the child. This can create problems for certain optimizing compilers if `<unistd.h>` is not included in the source calling **vfork()**.

NAME	vhangup – virtually “hangup” the current controlling terminal
SYNOPSIS	<pre>#include <unistd.h> void vhangup(void);</pre>
DESCRIPTION	The vhangup() function is used by the initialization process init(1M) (among others) to ensure that users are given “clean” terminals at login by revoking access of the previous users’ processes to the terminal. To effect this, vhangup() searches the system tables for references to the controlling terminal of the invoking process and revokes 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 or EIO). A SIGHUP (hangup signal) is sent to the process group of the controlling terminal.
SEE ALSO	init(1M)
BUGS	Access to the controlling terminal using <code>/dev/tty</code> is still possible. This call should be replaced by an automatic mechanism that takes place on process exit.

NAME	wait – wait for child process to stop or terminate
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/wait.h> pid_t wait(int *stat_loc);</pre>
DESCRIPTION	<p>The wait() function will suspend execution of the calling thread until status information for one of its terminated child processes is available, or until delivery of a signal whose action is either to execute a signal-catching function or to terminate the process. If more than one thread is suspended in wait() or waitpid(2) awaiting termination of the same process, exactly one thread will return the process status at the time of the target process termination. If status information is available prior to the call to wait(), return will be immediate.</p> <p>If wait() returns because the status of a child process is available, it returns the process ID of the child process. If the calling process specified a non-zero value for <i>stat_loc</i>, the status of the child process is stored in the location pointed to by <i>stat_loc</i>. That status may be evaluated with the macros described on the wstat(5) manual page.</p> <p>In the following, <i>status</i> is the object pointed to by <i>stat_loc</i>:</p> <ul style="list-style-type: none"> ■ If the child process stopped, the high order 8 bits of <i>status</i> will contain the number of the signal that caused the process to stop and the low order 8 bits will be set equal to WSTOPFLG. ■ If the child process terminated due to an _exit() call, the low order 8 bits of <i>status</i> will be 0 and the high order 8 bits will contain the low order 8 bits of the argument that the child process passed to _exit(); see exit(2). ■ If the child process terminated due to a signal, the high order 8 bits of <i>status</i> will be 0 and the low order 8 bits will contain the number of the signal that caused the termination. In addition, if WCOREFLG is set, a “core image” will have been produced; see signal(3C). <p>If the calling process has SA_NOCLDWAIT set or has SIGCHLD set to SIG_IGN, and the process has no unwaited children that were transformed into zombie processes, it will block until all of its children terminate, and wait() will fail and set errno to ECHILD.</p> <p>If a parent process terminates without waiting for its child processes to terminate, the parent process ID of each child process is set to 1, with the initialization process inheriting the child processes; see intro(2).</p>
RETURN VALUES	When wait() returns due to a terminated child process, the process ID of the child is returned to the calling process. Otherwise, -1 is returned and errno is set to indicate the error.

ERRORS

The **wait()** function will fail if:

- ECHILD** The calling process has no existing unwaited-for child processes.
- EINTR** The function was interrupted by a signal.

USAGE

Since **wait()** blocks on a stopped child, a calling process wishing to see the return results of such a call should use **waitid(2)** or **waitpid(2)** instead of **wait()**.

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

intro(2), **exec(2)**, **exit(2)**, **fork(2)**, **pause(2)**, **ptrace(2)**, **waitid(2)**, **waitpid(2)**, **signal(3C)**, **attributes(5)**, **signal(5)**, **wstat(5)**

NAME	waitid – wait for child process to change state												
SYNOPSIS	<pre>#include <wait.h> int waitid(idtype_t idtype, id_t id, siginfo_t *infp, int options);</pre>												
DESCRIPTION	<p>The waitid() function suspends the calling process until one of its child processes changes state. It records the current state of a child in the structure pointed to by <i>infp</i>. It returns immediately if a child process changed state prior to the call.</p> <p>The <i>idtype</i> and <i>id</i> arguments specify which children waitid() is to wait for, as follows:</p> <ul style="list-style-type: none"> ■ If <i>idtype</i> is <code>P_PID</code>, waitid() waits for the child with a process ID equal to <code>(pid_t)id</code>. ■ If <i>idtype</i> is <code>P_PGID</code>, waitid() waits for any child with a process group ID equal to <code>(pid_t)id</code>. ■ If <i>idtype</i> is <code>P_ALL</code>, waitid() waits for any child and <i>id</i> is ignored. <p>The <i>options</i> argument is used to specify which state changes waitid() is to wait for. It is formed by bitwise OR operation of any of the following flags:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><code>WCONTINUED</code></td> <td>Return the status for any child that was stopped and has been continued.</td> </tr> <tr> <td><code>WEXITED</code></td> <td>Wait for process(es) to exit.</td> </tr> <tr> <td><code>WNOHANG</code></td> <td>Return immediately.</td> </tr> <tr> <td><code>WNOWAIT</code></td> <td>Keep the process in a waitable state.</td> </tr> <tr> <td><code>WSTOPPED</code></td> <td>Wait for and return the process status of any child that has stopped upon receipt of a signal.</td> </tr> <tr> <td><code>WTRAPPED</code></td> <td>Wait for traced process(es) to become trapped or reach a breakpoint (see <code>ptrace(2)</code>).</td> </tr> </table> <p>The <i>infp</i> argument must point to a <code>siginfo_t</code> structure, as defined in <code>siginfo(5)</code>. If waitid() returns because a child process was found that satisfies the conditions indicated by the arguments <i>idtype</i> and <i>options</i>, then the structure pointed to by <i>infp</i> will be filled by the system with the status of the process. The <code>si_signo</code> member will always be equal to <code>SIGCHLD</code>.</p>	<code>WCONTINUED</code>	Return the status for any child that was stopped and has been continued.	<code>WEXITED</code>	Wait for process(es) to exit.	<code>WNOHANG</code>	Return immediately.	<code>WNOWAIT</code>	Keep the process in a waitable state.	<code>WSTOPPED</code>	Wait for and return the process status of any child that has stopped upon receipt of a signal.	<code>WTRAPPED</code>	Wait for traced process(es) to become trapped or reach a breakpoint (see <code>ptrace(2)</code>).
<code>WCONTINUED</code>	Return the status for any child that was stopped and has been continued.												
<code>WEXITED</code>	Wait for process(es) to exit.												
<code>WNOHANG</code>	Return immediately.												
<code>WNOWAIT</code>	Keep the process in a waitable state.												
<code>WSTOPPED</code>	Wait for and return the process status of any child that has stopped upon receipt of a signal.												
<code>WTRAPPED</code>	Wait for traced process(es) to become trapped or reach a breakpoint (see <code>ptrace(2)</code>).												
RETURN VALUES	<p>If waitid() returns due to a change of state of one of its children and <code>WNOHANG</code> was not used, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error. If <code>WNOHANG</code> was used, 0 can be returned (indicating no error); however, no children may have changed state if <code>info->si_pid</code> is 0.</p>												

ERRORS	The waitid() function will fail if: ECHILD The set of processes specified by <i>idtype</i> and <i>id</i> does not contain any unwaited processes. EFAULT The <i>infp</i> argument points to an illegal address. EINTR The waitid() function was interrupted due to the receipt of a signal by the calling process. EINVAL An invalid value was specified for <i>options</i> , or <i>idtype</i> and <i>id</i> specify an invalid set of processes.
USAGE	With <i>idtype</i> equal to <code>P_ALL</code> and <i>options</i> equal to <code>WEXITED WTRAPPED</code> , waitid() is equivalent to wait(2) .
SEE ALSO	intro(2) , exec(2) , exit(2) , fork(2) , pause(2) , ptrace(2) , sigaction(2) , wait(2) , signal(3C) , siginfo(5)

NAME	waitpid – wait for child process to change state
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/wait.h> pid_t waitpid(pid_t pid, int *stat_loc, int options);</pre>
DESCRIPTION	<p>The waitpid() function will suspend execution of the calling thread until status information for one of its terminated child processes is available, or until delivery of a signal whose action is either to execute a signal-catching function or to terminate the process. If more than one thread is suspended in waitpid() or wait(2) awaiting termination of the same process, exactly one thread will return the process status at the time of the target process termination. If status information is available prior to the call to waitpid(), return will be immediate.</p> <p>The <i>pid</i> argument specifies a set of child processes for which status is requested, as follows:</p> <ul style="list-style-type: none"> ■ If <i>pid</i> is equal to (pid_t)-1, status is requested for any child process. <ul style="list-style-type: none"> If <i>pid</i> is greater than (pid_t)0, it specifies the process ID of the child process for which status is requested. ■ If <i>pid</i> is equal to (pid_t)0 status is requested for any child process whose process group ID is equal to that of the calling process. ■ If <i>pid</i> is less than (pid_t)-1, status is requested for any child process whose process group ID is equal to the absolute value of <i>pid</i>. <p>If the calling process has SA_NOCLDWAIT set or has SIGCHLD set to SIG_IGN and the process has no unwaited children that were transformed into zombie processes, it will block until all of its children terminate, and waitpid() will fail and set <i>errno</i> to ECHILD.</p> <p>If waitpid() returns because the status of a child process is available, then that status may be evaluated with the macros defined by wstat(5). If the calling process had specified a non-zero value of <i>stat_loc</i>, the status of the child process will be stored in the location pointed to by <i>stat_loc</i>.</p> <p>The <i>options</i> argument is constructed from the bitwise inclusive OR of zero or more of the following flags, defined in the header <code><sys/wait.h></code>:</p> <p>WCONTINUED The status of any continued child process specified by <i>pid</i>, whose status has not been reported since it continued, is also reported to the calling process.</p> <p>WNOHANG waitpid() will not suspend execution of the calling process if status is not immediately available for one of the child processes specified by <i>pid</i>.</p>

WNOWAIT Keep the process whose status is returned in *stat_loc* in a waitable state. The process may be waited for again with identical results.

WUNTRACED The status of any child processes specified by *pid* that are stopped, and whose status has not yet been reported since they stopped, is also reported to the calling process.

RETURN VALUES

If **waitpid()** returns because the status of a child process is available, it returns a value equal to the process ID of the child process for which status is reported. If **waitpid()** returns due to the delivery of a signal to the calling process, `-1` is returned and `errno` is set to `EINTR`. If **waitpid()** 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*, then `0` is returned. Otherwise, `-1` is returned and `errno` is set to indicate the error.

ERRORS

The **waitpid()** function will fail if:

ECHILD The process or process group specified by *pid* does not exist or is not a child of the calling process or can never be in the states specified by *options*.

EINTR The **waitpid()** function was interrupted due to the receipt of a signal sent by the calling process.

EINVAL An invalid value was specified for *options*.

USAGE

With *options* equal to `0` and *pid* equal to `(pid_t)-1`, **waitpid()** is identical to **wait(2)**.

ATTRIBUTES

See **attributes(5)** for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Async-Signal-Safe

SEE ALSO

intro(2), **exec(2)**, **exit(2)**, **fork(2)**, **pause(2)**, **ptrace(2)**, **sigaction(2)**, **wait(2)**, **signal(3C)**, **attributes(5)**, **siginfo(5)**, **wstat(5)**

NAME	write, pwrite, writev – write on a file
SYNOPSIS	<pre>#include <unistd.h> ssize_t write(int fildes, const void * buf, size_t nbytes); ssize_t pwrite(int fildes, const void * buf, size_t nbytes, off_t offset); #include <sys/uio.h> ssize_t writev(int fildes, const struct iovec * iov, int iovcnt);</pre>
DESCRIPTION	<p>The write() function attempts to write <i>nbyte</i> bytes from the buffer pointed to by <i>buf</i> to the file associated with the open file descriptor, <i>fildes</i> .</p> <p>If <i>nbyte</i> is 0, write() will return 0 and have no other results if the file is a regular file; otherwise, the results are unspecified.</p> <p>On a regular file or other file capable of seeking, the actual writing of data proceeds from the position in the file indicated by the file offset associated with <i>fildes</i> . Before successful return from write() , the file offset is incremented by the number of bytes actually written. On a regular file, if this incremented file offset is greater than the length of the file, the length of the file will be set to this file offset.</p> <p>If the <code>O_SYNC</code> flag of the file status flags is set and <i>fildes</i> refers to a regular file, a successful write() does not return until the data is delivered to the underlying hardware.</p> <p>If <i>fildes</i> refers to a socket, write() is equivalent to <code>send(3N)</code> with no flags set.</p> <p>On a file not capable of seeking, writing always takes place starting at the current position. The value of a file offset associated with such a device is undefined.</p> <p>If the <code>O_APPEND</code> flag of the file status flags is set, the file offset will be set to the end of the file prior to each write and no intervening file modification operation will occur between changing the file offset and the write operation.</p> <p>For regular files, no data transfer will occur past the offset maximum established in the open file description with <i>fildes</i> .</p> <p>A write() to a regular file is blocked if mandatory file/record locking is set (see <code>chmod(2)</code>), and there is a record lock owned by another process on the segment of the file to be written:</p> <ul style="list-style-type: none"> ■ If <code>O_NDELAY</code> or <code>O_NONBLOCK</code> is set, write() returns <code>-1</code> and sets <code>errno</code> to <code>EAGAIN</code> . ■ If <code>O_NDELAY</code> and <code>O_NONBLOCK</code> are clear, write() sleeps until all blocking locks are removed or the write() is terminated by a signal.

If a **write()** requests that more bytes be written than there is room for—for example, if the write would exceed the process file size limit (see **getrlimit(2)** and **ulimit(2)**), the system file size limit, or the free space on the device—only as many bytes as there is room for will be written. For example, suppose there is space for 20 bytes more in a file before reaching a limit. A **write()** of 512-bytes returns 20. The next **write()** of a non-zero number of bytes gives a failure return (except as noted for pipes and FIFO below).

If **write()** is interrupted by a signal before it writes any data, it will return `-1` with `errno` set to `EINTR`.

If **write()** is interrupted by a signal after it successfully writes some data, it will return the number of bytes written.

If the value of `nbyte` is greater than `SSIZE_MAX`, the result is implementation-dependent.

After a **write()** to a regular file has successfully returned:

- Any successful **read(2)** from each byte position in the file that was modified by that write will return the data specified by the **write()** for that position until such byte positions are again modified.
- Any subsequent successful **write()** to the same byte position in the file will overwrite that file data.

Write requests to a pipe or FIFO are handled the same as a regular file with the following exceptions:

- There is no file offset associated with a pipe, hence each write request appends to the end of the pipe.
- Write requests of `{PIPE_BUF}` bytes or less are guaranteed not to be interleaved with data from other processes doing writes on the same pipe. Writes of greater than `{PIPE_BUF}` bytes may have data interleaved, on arbitrary boundaries, with writes by other processes, whether or not the `O_NONBLOCK` or `O_NDELAY` flags are set.
- If `O_NONBLOCK` and `O_NDELAY` are clear, a write request may cause the process to block, but on normal completion it returns `nbyte`.
- If `O_NONBLOCK` and `O_NDELAY` are set, **write()** does not block the process. If a **write()** request for `PIPE_BUF` or fewer bytes succeeds completely **write()** returns `nbyte`. Otherwise, if `O_NONBLOCK` is set, it returns `-1` and sets `errno` to `EAGAIN` or if `O_NDELAY` is set, it returns `0`. A **write()** request for greater than `{PIPE_BUF}` bytes transfers what it can and returns the number of bytes written or it transfers no data and, if `O_NONBLOCK` is set, returns `-1` with `errno` set to `EAGAIN` or if `O_NDELAY` is set, it returns `0`. Finally, if a 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.

When attempting to write to a file descriptor (other than a pipe, a FIFO, a socket, or a STREAM) that supports nonblocking writes and cannot accept the data immediately:

- If `O_NONBLOCK` and `O_NDELAY` are clear, **write()** blocks until the data can be accepted.
- If `O_NONBLOCK` or `O_NDELAY` is set, **write()** does not block the process. If some data can be written without blocking the process, **write()** writes what it can and returns the number of bytes written. Otherwise, if `O_NONBLOCK` is set, it returns `-1` and sets `errno` to `EAGAIN` or if `O_NDELAY` is set, it returns `0`.

Upon successful completion, where *nbyte* is greater than 0, **write()** will mark for update the `st_ctime` and `st_mtime` fields of the file, and if the file is a regular file, the `S_ISUID` and `S_ISGID` bits of the file mode may be cleared.

For STREAMS files (see [intro\(2\)](#) and [streamio\(7I\)](#)), the operation of **write()** is determined by the values of the minimum and maximum *nbyte* range ("packet size") accepted by the STREAM. These values are contained in the topmost STREAM module, and can not be set or tested from user level. If *nbyte* falls within the packet size range, *nbyte* bytes are written. If *nbyte* does not fall within the range and the minimum packet size value is zero, **write()** breaks the buffer into maximum packet size segments prior to sending the data downstream (the last segment may be smaller than the maximum packet size). If *nbyte* does not fall within the range and the minimum value is non-zero, **write()** fails and sets `errno` to `ERANGE`. Writing a zero-length buffer (*nbyte* is zero) to a STREAMS device sends a zero length message with zero returned. However, writing a zero-length buffer to a pipe or FIFO sends no message and zero is returned. The user program may issue the `I_SWROPT ioctl(2)` to enable zero-length messages to be sent across the pipe or FIFO (see [streamio\(7I\)](#)).

When writing to a STREAM, data messages are created with a priority band of zero. When writing to a socket or to a STREAM that is not a pipe or a FIFO:

- If `O_NDELAY` and `O_NONBLOCK` are not set, and the STREAM cannot accept data (the STREAM write queue is full due to internal flow control conditions), **write()** blocks until data can be accepted.
- If `O_NDELAY` or `O_NONBLOCK` is set and the STREAM cannot accept data, **write()** returns `-1` and sets `errno` to `EAGAIN`.
- If `O_NDELAY` or `O_NONBLOCK` is set and part of the buffer has already been written when a condition occurs in which the STREAM cannot accept additional data, **write()** terminates and returns the number of bytes written.

The **write()** and **writev()** functions will fail if the STREAM head had processed an asynchronous error before the call. In this case, the value of `errno` does not reflect the result of **write()** or **writev()** but reflects the prior error.

pwrite() The **pwrite()** function performs the same action as **write()**, except that it writes into a given position without changing the file pointer. The first three arguments to **pwrite()** are the same as **write()** with the addition of a fourth argument *offset* for the desired position inside the file.

writev() The **writev()** function performs the same action as **write()**, but gathers the output data from the *iovcnt* buffers specified by the members of the *iov* array: *iov* [0], *iov* [1], ..., *iov* [*iovcnt* - 1]. The *iovcnt* buffer is valid if greater than 0 and less than or equal to {`IOV_MAX`}. See **intro(2)** for a definition of {`IOV_MAX`}.

The `iovec` structure contains the following members:

```
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. The **writev()** function always writes all data from an area before proceeding to the next.

If *fildev* refers to a regular file and all of the `iov_len` members in the array pointed to by *iov* are 0, **writev()** will return 0 and have no other effect. For other file types, the behavior is unspecified.

If the sum of the `iov_len` values is greater than `SSIZE_MAX`, the operation fails and no data is transferred.

RETURN VALUES

Upon successful completion, **write()** returns the number of bytes actually written to the file associated with *fildev*. This number is never greater than *nbyte*. Otherwise, `-1` is returned, the file-pointer remains unchanged, and `errno` is set to indicate the error.

Upon successful completion, **writev()** returns the number of bytes actually written. Otherwise, it returns `-1`, the file-pointer remains unchanged, and `errno` is set to indicate an error.

ERRORS

The **write()**, **pwrite()**, and **writev()** functions will fail if:

EAGAIN Mandatory file/record locking is set, `O_NDELAY` or `O_NONBLOCK` is set, and there is a blocking record lock; total amount of system memory available when reading using

	raw I/O is temporarily insufficient; an attempt is made to write to a STREAM that can not accept data with the <code>O_NDELAY</code> or <code>O_NONBLOCK</code> flag set; or a write to a pipe or FIFO of <code>PIPE_BUF</code> bytes or less is requested and less than <i>nbytes</i> of free space is available.
EBADF	The <i>fildev</i> argument is not a valid file descriptor open for writing.
EDEADLK	The write was going to go to sleep and cause a deadlock situation to occur.
EDQUOT	The user's quota of disk blocks on the file system containing the file has been exhausted.
EFAULT	The <i>buf</i> argument points to an illegal address.
EFBIG	An attempt is made to write a file that exceeds the process's file size limit or the maximum file size (see <code>getrlimit(2)</code> and <code>ulimit(2)</code>).
EFBIG	The file is a regular file, <i>nbyte</i> is greater than 0, and the starting position is greater than or equal to the offset maximum established in the file description associated with <i>fildev</i> .
EINTR	A signal was caught during the write operation and no data was transferred.
EIO	The process is in the background and is attempting to write to its controlling terminal whose <code>TOSTOP</code> flag is set, or the process is neither ignoring nor blocking <code>SIGTTOU</code> signals and the process group of the process is orphaned.
ENOLCK	Enforced record locking was enabled and <code>{LOCK_MAX}</code> regions are already locked in the system, or the system record lock table was full and the write could not go to sleep until the blocking record lock was removed.
ENOLINK	The <i>fildev</i> argument is on a remote machine and the link to that machine is no longer active.
ENOSPC	During a write to an ordinary file, there is no free space left on the device.
ENOSR	An attempt is made to write to a STREAMS with insufficient STREAMS memory resources available in the system.

- ENXIO** A hangup occurred on the STREAM being written to.
- EPIPE** An attempt is made to write to a pipe or a FIFO that is not open for reading by any process, or that has only one end open (or to a file descriptor created by `socket(3N)`, using type `SOCK_STREAM` that is no longer connected to a peer endpoint). A `SIGPIPE` signal will also be sent to the process. The process dies unless special provisions were taken to catch or ignore the signal.
- ERANGE** The transfer request size was outside the range supported by the STREAMS file associated with *fildev*.
- The `pwrite()` function fails and the file pointer remains unchanged if:
- ESPIPE** The *fildev* argument is associated with a pipe or FIFO.
- The `writev()` function will fail if:
- EINVAL** The sum of the `iov_len` values in the *iov* array would overflow an `ssize_t`.
- The `write()` and `writev()` functions may fail if:
- EINVAL** The STREAM or multiplexer referenced by *fildev* is linked (directly or indirectly) downstream from a multiplexer.
- ENXIO** A request was made of a non-existent device, or the request was outside the capabilities of the device.
- ENXIO** A hangup occurred on the STREAM being written to. A write to a STREAMS file may 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.
- The `writev()` function may fail if:
- EINVAL** The *iovcnt* argument was less than or equal to 0 or greater than `{ IOV_MAX }`; one of the `iov_len` values in the *iov* array was negative; or the sum of the `iov_len` values in the *iov* array overflowed an `int`.

USAGE

The `pwrite()` function has a transitional interface for 64-bit file offsets. See `1f64(5)`.

ATTRIBUTES

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	<code>write()</code> is Async-Signal-Safe

SEE ALSO

`Intro(2)`, `chmod(2)`, `creat(2)`, `dup(2)`, `fcntl(2)`, `getrlimit(2)`,
`ioctl(2)`, `lseek(2)`, `open(2)`, `pipe(2)`, `ulimit(2)`, `send(3N)`,
`socket(3N)`, `attributes(5)`, `lf64(5)`, `streamio(7I)`

NAME	yield – yield execution to another lightweight process
SYNOPSIS	<pre>#include <unistd.h> void yield(void);</pre>
DESCRIPTION	The yield() function causes the current lightweight process to yield its execution in favor of another lightweight process with the same or greater priority.
SEE ALSO	thr_yield(3T)

Index

A

access — determine accessibility of a file, 36
access permission mode of file
 change — `chmod`, 58
accounting
 enable or disable process accounting —
 `acct`, 38
`acct` — enable or disable process
 accounting, 38
`acl` — get or set a file's Access Control List
 (ACL), 40
`adjtime` — correct the time to allow
 synchronization of the system
 clock, 42
adjust local clock parameters —
 `ntp_adjtime`, 175
alarm — set a process alarm clock, 44
audit — write an audit record, 46
auditon — manipulate auditing, 47
auditsvc() function, 52

B

bind LWPs to a processor —
 `processor_bind`, 209
bind LWPs to a set of processors —
 `pset_bind`, 214
`brk` — change the amount of space allocated
 for the calling process's data
 segment, 54

C

`chdir` — change working directory, 56

child processes

allows a parent process to control the
 execution of a child process
 — `ptrace`, 219
get time — `times`, 307
wait for child process to change state —
 `waitid`, 328, 330
wait for child process to stop or terminate
 — `wait`, 326
`chmod` — change access permission mode of
 file, 58
`chown` — change owner and group of a file, 62
`chroot` — change root directory, 65
clock
 get local clock values — `ntp_gettime`, 176
continue or suspend LWP execution
 — `_lwp_continue`, 143
 — `_lwp_suspend`, 143
CPU-use
 process execution time profile —
 `profil`, 212
`creat` — create a new file or rewrite an existing
 one, 69
create a new process — `fork`, 90
 `fork1`, 90
create a new light-weight process —
 `_lwp_create`, 132
create session and set process group ID —
 `setsid`, 252

D

deliver process signals to specific LWPs

- `_lwp_sigredirect`, 267
- `_signotifywait`, 267
- determine accessibility of a file — `access`, 36
- devices
 - I/O control functions — `ioctl`, 120
- directories
 - change working directory — `chdir`, 56
 - create a new one — `mknod`, 152
 - get configurable pathname variables — `pathconf`, 94
 - make a new one — `mkdir`, 150
 - read directory entries and put in a file system independent format — `getdents`, 103
 - remove — `rmdir`, 237
- `dup` — duplicate an open file descriptor, 72

E

- effective group ID
 - set — `setregid()`, 250
- effective user ID
 - set — `setreuid()`, 251
- `exec` — execute a file, 73
- `execl` — execute a file, 73
- `execle` — execute a file, 73
- `execlp` — execute a file, 73
- `execv` — execute a file, 73
- `execve` — execute a file, 73
- `execvp` — execute a file, 73
- `exit` — terminate process, 78
- `_exit` — terminate process, 78

F

- `facl` — get or set a file's Access Control List (ACL), 40
- `fchdir` — change working directory, 56
- `fchmod` — change access permission mode of file, 58
- `fchown` — change owner and group of a file, 62
- `fcntl` — file control, 81
- file control — `fcntl`, 81
- file descriptor
 - duplicate an open one — `dup`, 72
- file pointer, read/write
 - move — `lseek`, 126, 127

- file status
 - get — `stat`, `lstat`, `fstat`, 286
- file system
 - get information — `statvfs`, `fstatvfs`, 290
 - get statistics — `ustat`, 318
 - make a symbolic link to a file — `symlink`, 298
 - remove link — `unlink`, 316
 - returns information about the file system types configured in the system — `sysfs`, 301
 - unmount — `umount`, 314
 - update super block — `sync`, 300

files

- change access permission mode of file — `chmod`, 58, 62
- change the name of a file — `rename`, 232
- create a new file or rewrite an existing one — `creat`, 69
- execute — `exec`, 73
- get configurable pathname variables — `pathconf`, 94
- link to a file — `link`, 124
- move read/write file pointer — `lseek`, 126, 127
- set file access and modification times — `utime`, 319
- `fork` — create a new process, 90
 - spawn new process in a virtual memory efficient way — `vfork`, 323
- `fork1` — create a new process, 90
- `fpathconf` — get configurable pathname variables, 94
- `fstat` — get status on open file known by file descriptor, 286
- `fstatvfs` — get file system information, 290

G

- get and set process limits — `ulimit`, 311
- get information about a processor set — `pset_info`, 218
- get LWP identifier — `_lwp_self`, 139
- get or change processor operational status — `p_online`, 193
- get or set a file's Access Control List (ACL) — `acl`, 40

- fac, 40
- get process group ID of session leader —
 - getsid, 118
- getaudit get process audit information, 97
- getaudit — get user audit identity, 99
- getdents — read directory entries and put in a
 - file system independent
 - format, 103
- getegid — get effective group ID, 119
- geteuid — get effective user ID, 119
- getgid — get real group ID, 119
- getgroups — get supplementary group access
 - list IDs, 105
- getitimer — get value of interval timer, 106
- getmsg — get next message off a stream, 110
- getpgid — get process group IDs, 113
- getpgrp — get process group IDs, 113
- getpid — get process IDs, 113
- getpmsg — get next message off a stream, 110
- getppid — get parent process IDs, 113
- getrlimit — control maximum system resource
 - consumption, 114
- getsid — get process group ID of session
 - leader, 118
- getuid — get real user ID, 119
- group ID
 - set real and effective — setregid(), 250
- group IDs
 - get — getgid, getegid, 119
 - set — setgid, 253
 - supplementary group access list IDs —
 - getgroups, setgroups, 105

H

- halt system
 - uadmin, 309
- hangup signal
 - the current controlling terminal —
 - vhangup, 325

I

- I/O
 - audit — audit, 46
 - multiplexing — poll, 190
- initialize an LWP context —
 - _lwp_makecontext, 137

- interprocess communication
 - pipe, 189
- interval timer
 - get or set value of interval timer —
 - getitimer, setitimer, 106
- ioctl — control device, 120

K

- kill — send a signal to a process or a group of
 - processes, 122

L

- lchown — change owner and group of a
 - file, 62
- link — link to a file, 124
 - remove — unlink, 316
- link, symbolic
 - make one to a file — symlink, 298
- lseek — move extended read/write file
 - pointer, 126
- lseek — move read/write file pointer, 127
- lstat — get status on symbolic link file, 286
- LWP
 - scheduler control — priocntl, 196
 - _lwp_cond_broadcast — signal a condition
 - variable, 129
 - _lwp_cond_signal — signal a condition
 - variable, 129
 - _lwp_cond_timedwait — wait on a condition
 - variable, 130
 - _lwp_cond_wait — wait on a condition
 - variable, 130
 - _lwp_continue — continue or suspend LWP
 - execution, 143
 - _lwp_create — create a new light-weight
 - process, 132
 - _lwp_exit — terminate the calling LWP, 134
 - _lwp_getprivate — set/get LWP specific
 - storage, 142
 - _lwp_info — return the time-accounting
 - information of a single
 - LWP, 135
 - _lwp_kill — send a signal to a LWP, 136
 - _lwp_makecontext — initialize an LWP
 - context, 137

- `_lwp_mutex_lock` — mutual exclusion, 138
- `_lwp_mutex_trylock` — mutual exclusion, 138
- `_lwp_mutex_unlock` — mutual exclusion, 138
- `_lwp_self` — get LWP identifier, 139
- `_lwp_sema_init` — semaphore operations, 140
- `_lwp_sema_post` — semaphore operations, 140
- `_lwp_sema_trywait` — semaphore operations, 140
- `_lwp_sema_wait` — semaphore operations, 140
- `_lwp_setprivate` — set/get LWP specific storage, 142
- `_lwp_sigredirect` — deliver process signals to specific LWPs, 267
- `_lwp_suspend` — continue or suspend LWP execution, 143
- `_lwp_wait` — wait for a LWP to terminate, 144

M

- make a directory, or a special or ordinary file
 - `mknod`, 152
- manage sets of processors
 - `pset_assign`, 216
 - `pset_create`, 216
 - `pset_destroy`, 216
- manipulate auditing — `auditon`, 47
- masks
 - set and get file creation mask — `umask`, 313
- `memcntl` — memory management control, 145
- memory
 - management control — `memcntl`, 145
- memory management
 - change the amount of space allocated for the calling process's data segment — `brk`, `sbrk`, 54
- memory mapping
 - set protection — `mprotect`, 163
- memory pages
 - determine residency — `mincore`, 149
 - map — `mmap`, 155
 - unmap — `munmap`, 173
- memory, shared
 - control operations — `shmctl`, 255
 - get segment identifier — `shmget`, 257
 - operations — `shmop`, 259
- message control operations
 - `msgctl`, 165

- message queue
 - get — `msgget`, 167
- message receive operation — `msgrcv`, 169
- message send operation — `msgsnd`, 171
- messages
 - send a message on a stream — `putmsg`, 221
- `mincore` — determine residency of memory pages, 149
- `mkdir` — make a directory, 150
- `mknod` — make a directory, or a special or ordinary file, 152
- `mmap` — map pages of memory, 155
- `mount` — mount a file system, 161
- mount a file system — `mount`, 161
- `mprotect` — set protection of memory mapping, 163
- `msgctl` — message control operations, 165
- `msgget` — get message queue, 167
- `msgrcv` — message receive operation, 169
- `msgsnd` — message send operation, 171
- `munmap` — unmap pages of memory, 173
- mutual exclusion
 - `_lwp_mutex_lock`, 138
 - `_lwp_mutex_trylock`, 138
 - `_lwp_mutex_unlock`, 138

N

- `nice` — change priority of a time-sharing process, 174
- `ntp_adjtime` — adjust local clock parameters, 175
- `ntp_gettime` — get local clock values, 176

O

- `open` — open a file, 177
- open a file — `open`, 177
- operating system
 - get name of current one — `uname`, 315
- owner of file
 - change — `chown`, 62

P

- `pathconf` — get configurable pathname variables, 94

pathname
 get configurable variables — pathconf, 94
 pause — suspend process until signal, 186
 pipe — create an interprocess channel, 189
 poll — input/output multiplexing, 190
 pread — read from file, 224
 prctl — process scheduler control, 196
 prctlset — generalized process scheduler control, 207
 process accounting
 enable or disable — acct, 38
 process alarm clock
 set — alarm, 44
 process audit information
 get process audit information —
 getaudit, 97
 set process audit information —
 setaudit, 97
 process group ID
 set — setpgid, 247, 249
 process scheduler
 control — prctl, 196
 generalized control — prctlset, 207
 process statistics
 process execution time profile —
 profil, 212
 process, time-sharing
 change priority — nice, 174
 processes
 allows a parent process to control the
 execution of a child process
 — ptrace, 219
 change priority of a time-sharing process
 — nice, 174
 create a new one — fork, 90
 create an interprocess channel — pipe, 189
 execute a file — exec, 73
 execution time profile — profil, 212
 generalized scheduler control —
 prctlset, 207
 get or set value of interval timer —
 getitimer, setitimer, 106, 110,
 113, 119
 get time — times, 307
 read from file — read, 224
 read directory entries and put in a file
 system independent format
 — getdents, 103
 send a signal to a process or a group of
 processes — kill, 122
 set a process alarm clock — alarm, 44
 set and get file creation mask —
 umask, 313
 set process group ID — setpgid, 247, 249
 spawn new process in a virtual memory
 efficient way — vfork, 323
 supplementary group access list IDs —
 getgroups, setgroups, 105
 suspend process until signal — pause, 186
 the current controlling terminal —
 vhangup, 325
 wait for child process to change state —
 waitid, 328, 330
 wait for child process to stop or terminate
 — wait, 326
 processes and protection
 — setregid(), 250
 — setreuid(), 251
 processor_bind — bind LWPs to a
 processor, 209
 processor_info — determine type and status of
 a processor, 211
 profil — process execution time profile, 212
 profiling utilities
 execution time profile — profil, 212
 pset_assign — manage sets of processors, 216
 pset_bind — bind LWPs to a set of
 processors, 214
 pset_create — manage sets of processors, 216
 pset_destroy — manage sets of processors, 216
 pset_info — get information about a processor
 set, 218
 ptrace — allows a parent process to control
 the execution of a child
 process, 219
 putmsg — send a message on a stream, 221
 putpmsg — send a message on a stream, 221
 pwrite — write on a file, 332
 p_online — get or change processor
 operational status, 193

R

read from file — read, 224
 pread, 224

- readv, 224
- read the contents of a symbolic link —
 readlink, 230
- read/write file pointer
 - move — lseek, 126, 127
- readlink — read the contents of a symbolic
 link, 230
- read — read from file, 224
- real group ID
 - set — setregid(), 250
- real user ID
 - set — setreuid(), 251
- reboot system
 - uadmin, 309
- remount root file system
 - uadmin, 309
- rename — change the name of a file, 232
- resolve all symbolic links of a path name —
 resolvepath, 235
- resolvepath — resolve all symbolic links of a
 path name, 235
- rmdir — remove a directory, 237
- root directory
 - change — chroot, 65

S

- sbrk — change the amount of space allocated
 for the calling process's data
 segment, 54
- semaphore operations
 - _lwp_sema_init, 140
 - _lwp_sema_post, 140
 - _lwp_sema_trywait, 140
 - _lwp_sema_wait, 140
- semaphores
 - control operations — semctl, 239
 - get a set — semget, 242
 - operations — semop, 244
- semctl — semaphore control operations, 239
- semget — get set of semaphores, 242
- semop — semaphore operations, 244
- send a signal to a LWP — _lwp_kill, 136
- set file access and modification times —
 utimes, 321
- set/get LWP specific storage
 - _lwp_getprivate, 142
 - _lwp_setprivate, 142

- setaudit set process audit information, 97
- setaudit — set user audit identity, 99
- setegid — set effective group ID, 253
- seteuid — set effective user ID, 253
- setgid — set group ID, 253
- setgroups — set supplementary group access
 list IDs, 105
- setitimer — set value of interval timer, 106
- setpgid — set process group ID, 247
- setpgrp — set process group ID, 249
- setregid() — set real and effective group
 ID, 250
- setreuid() — set real and effective user IDs, 251
- setrlimit — control maximum system resource
 consumption, 114
- setsid — create session and set process group
 ID, 252
- setuid — set user ID, 253
- shared memory
 - control operations — shmctl, 255
 - get segment identifier — shmget, 257
 - operations — shmop, 259
- shmctl — shared memory control
 operations, 255
- shmget — get shared memory segment
 identifier, 257
- shmop — shared memory operations, 259
- shutdown
 - uadmin, 309
- sigaction — detailed signal management, 262
- sigaltstack — set or get signal alternate stack
 context, 265
- signal a condition variable
 - _lwp_cond_broadcast, 129
 - _lwp_cond_signal, 129
- signal alternate stack
 - set or get context — sigaltstack, 265
- signal management
 - detailed — sigaction, 262
- signal mask
 - change and/or examine —
 sigprocmask, 270
 - install, and suspend process until signal
 — sigsuspend, 274
- signals
 - examine blocked and pending ones —
 sigpending, 269

- `_signotifywait` — deliver process signals to specific LWPs, 267
- `sigpending` — examine signals that are blocked and pending, 269
- `sigprocmask` — change and/or examine calling process's signal mask, 270
- `sigsend` — send a signal to a process or a group of processes, 272
- `sigsendset` — provides an alternate interface to `sigsend` for sending signals to sets of processes, 272
- `sigsuspend` — install a signal mask and suspend process until signal, 274
- `sigwait()` — wait until a signal is posted, 276
- install a SPARC V9 user trap handler, 280
- special files
 - create a new one — `mknod`, 152
- `stat` — get file status, 286
- statistics
 - get for mounted file system — `ustat`, 318
- `statvfs` — get file system information, 290
- `stime` — set system time and date, 293
- STREAMS
 - get next message off a stream — `getmsg`, 110
 - I/O control functions — `ioctl`, 120
 - send a message on a stream — `putmsg`, 221
- super block
 - update — `sync`, 300
- swap space
 - manage — `swapctl`, 294
- `swapctl` — manage swap space, 294
- symbolic link
 - make one to a file — `symlink`, 298
- `symlink` — make a symbolic link to a file, 298
- `sync` — update super block, 300
- `sysinfo` — get and set system information strings, 302
- system administration
 - administrative control — `uadmin`, 309
- system clock
 - synchronization — `adjtime`, 42
- system information
 - get and set strings — `sysinfo`, 302
- system operation
 - update super block — `sync`, 300

- system resources
 - control maximum system resource consumption — `getrlimit`, `setrlimit`, 114

T

- terminate process
 - `exit`, 78
 - `_exit`, 78
- terminate the calling LWP — `_lwp_exit`, 134
- time — get time, 306
 - correct the time to allow synchronization of the system clock — `adjtime`, 42
 - set system time and date — `stime`, 293
- time-accounting
 - single LWP — `_lwp_info`, 135
- times — get process and child process times, 307

U

- `ulimit` — get and set process limits, 311
- `umask` — set and get file creation mask, 313
- `umount` — unmount a file system, 314
- `uname` — get name of current operating system, 315
- `unlink` — remove directory entry, 316
- unmount a file system — `umount`, 314
- user audit identity
 - get user audit identity — `getaudit`, 99
 - set user audit identity — `setaudit`, 99
- user ID
 - set real and effective — `setreuid()`, 251
- user IDs
 - get — `getuid`, `geteuid`, 119
 - set — `setuid`, 253
- `utime` — set file access and modification times, 319
- `utimes` — set file access and modification times, 321

V

- `vfork` — spawn new process in a virtual memory efficient way, 323

vhangup — the current controlling terminal, 325

W

wait — wait for child process to stop or terminate, 326
wait for a LWP to terminate — `_lwp_wait`, 144
wait on a condition variable
— `_lwp_cond_timedwait`, 130
— `_lwp_cond_wait`, 130
waitid — wait for child process to change state, 328
waitpid — wait for child process to change state, 330

write on a file
— write, 332
— write, 332
— write, 332
write — write on a file, 332

Y

yield — yield execution to another lightweight process, 339
yield execution to another lightweight process
— yield, 339