



man pages section 3: Basic Library Functions

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 816-5168-10
January 2005

Copyright 2005 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 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, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks or registered trademarks 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.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

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 2005 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. 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 États-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, de Sun Microsystems, Inc. aux États-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 États-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.



050105@10536



Contents

Preface 17

Basic Library Functions 23

a64l(3C) 24
abort(3C) 25
abs(3C) 26
addsev(3C) 27
addseverity(3C) 28
assert(3C) 30
atexit(3C) 31
atomic_ops(3C) 32
attropen(3C) 34
basename(3C) 35
bsdmalloc(3MALLOC) 36
bsd_signal(3C) 38
bsearch(3C) 39
bstring(3C) 41
btowc(3C) 42
catgets(3C) 43
catopen(3C) 44
cfgetispeed(3C) 47
cfsetispeed(3C) 48
clock(3C) 49
closedir(3C) 50
closefrom(3C) 51
cond_init(3C) 53

confstr(3C) 58
crypt(3C) 64
crypt_genhash_impl(3C) 66
crypt_gensalt(3C) 67
crypt_gensalt_impl(3C) 68
cset(3C) 69
ctermid(3C) 70
ctime(3C) 71
ctype(3C) 76
cuserid(3C) 79
dbm(3UCB) 80
decimal_to_floating(3C) 82
difftime(3C) 83
directio(3C) 84
dirname(3C) 86
div(3C) 88
dladdr(3C) 89
dlclose(3C) 91
dlldump(3C) 93
dlerror(3C) 99
dlnfo(3C) 100
dlopen(3C) 104
dlsym(3C) 109
drand48(3C) 112
dup2(3C) 114
econvert(3C) 115
ecvt(3C) 117
encrypt(3C) 119
end(3C) 120
euclen(3C) 121
exit(3C) 122
fattach(3C) 123
__fbufsize(3C) 125
fclose(3C) 127
fdetach(3C) 129
fdopen(3C) 131
ferror(3C) 133
fflush(3C) 134

ffs(3C) 136
fgetc(3C) 137
fgetpos(3C) 140
fgetwc(3C) 141
floating_to_decimal(3C) 143
flock(3UCB) 145
flockfile(3C) 147
fmtmsg(3C) 149
fnmatch(3C) 154
fopen(3C) 156
fopen(3UCB) 159
fpgetround(3C) 161
fputc(3C) 163
fputwc(3C) 166
fputws(3C) 168
fread(3C) 169
freopen(3C) 171
fseek(3C) 174
fsetpos(3C) 177
fsync(3C) 178
ftell(3C) 180
ftime(3C) 181
ftok(3C) 182
ftw(3C) 184
fwide(3C) 189
fwprintf(3C) 190
fwrite(3C) 199
fwscanf(3C) 200
getcpuid(3C) 207
getcwd(3C) 208
getdate(3C) 210
getdtablesize(3C) 216
getenv(3C) 217
getexecname(3C) 218
getgrnam(3C) 219
gethostid(3C) 223
gethostname(3C) 224
gethrtime(3C) 225

getloadavg(3C) 226
getlogin(3C) 227
getmntent(3C) 229
getnetgrent(3C) 231
getopt(3C) 233
getopt_long(3C) 238
getpagesize(3C) 246
getpagesizes(3C) 247
getpass(3C) 248
getpeerucred(3C) 249
getpriority(3C) 251
getpw(3C) 254
getpwnam(3C) 255
getrusage(3C) 260
gets(3C) 263
getspnam(3C) 264
getsubopt(3C) 268
gettext(3C) 272
gettimeofday(3C) 276
gettimeofday(3UCB) 278
gettxt(3C) 279
getusershell(3C) 281
getutent(3C) 282
getutxent(3C) 285
getvfsent(3C) 289
getwc(3C) 291
getwchar(3C) 292
getwd(3C) 293
getwidth(3C) 294
getws(3C) 295
getzoneid(3C) 297
glob(3C) 298
grantpt(3C) 302
hsearch(3C) 303
iconv(3C) 306
iconv_close(3C) 312
iconv_open(3C) 313
imaxabs(3C) 315

imaxdiv(3C) 316
index(3C) 317
initgroups(3C) 318
insque(3C) 319
isaexec(3C) 320
isastream(3C) 321
isatty(3C) 322
isnand(3C) 323
iswalpha(3C) 325
iswctype(3C) 328
killpg(3C) 330
lckpwwdf(3C) 331
lfmt(3C) 332
localeconv(3C) 336
lockf(3C) 342
_longjmp(3C) 345
lsearch(3C) 346
madvise(3C) 348
makecontext(3C) 351
makedev(3C) 354
malloc(3C) 355
malloc(3MALLOC) 358
mapmalloc(3MALLOC) 361
mblen(3C) 363
mbrlen(3C) 364
mbrtowc(3C) 366
mbsinit(3C) 368
mbsrtowcs(3C) 369
mbstowcs(3C) 371
mbtowc(3C) 372
mctl(3UCB) 373
memory(3C) 375
mkfifo(3C) 377
mkstemp(3C) 379
mktemp(3C) 380
mktime(3C) 381
mlock(3C) 384
mlockall(3C) 386

monitor(3C) 388
msync(3C) 390
mtmalloc(3MALLOC) 392
mutex_init(3C) 395
ndbm(3C) 407
nice(3UCB) 411
nlist(3UCB) 412
nl_langinfo(3C) 413
offsetof(3C) 414
opendir(3C) 415
perror(3C) 417
pfmt(3C) 418
plock(3C) 421
popen(3C) 422
port_alert(3C) 425
port_associate(3C) 427
port_create(3C) 431
port_get(3C) 433
port_send(3C) 437
posix_openpt(3C) 440
posix_spawn(3C) 442
posix_spawnattr_destroy(3C) 447
posix_spawnattr_getflags(3C) 449
posix_spawnattr_getpgroup(3C) 451
posix_spawnattr_getschedparam(3C) 452
posix_spawnattr_getschedpolicy(3C) 453
posix_spawnattr_getsigdefault(3C) 454
posix_spawnattr_getsigmask(3C) 455
posix_spawn_file_actions_addclose(3C) 456
posix_spawn_file_actions_adddup2(3C) 458
posix_spawn_file_actions_destroy(3C) 459
printf(3C) 460
printf(3UCB) 471
priv_addset(3C) 475
priv_set(3C) 478
priv_str_to_set(3C) 480
pset_getloadavg(3C) 483
psignal(3C) 484

psignal(3UCB) 485
pthread_atfork(3C) 486
pthread_attr_getdetachstate(3C) 488
pthread_attr_getguardsize(3C) 489
pthread_attr_getinheritsched(3C) 491
pthread_attr_getschedparam(3C) 493
pthread_attr_getschedpolicy(3C) 494
pthread_attr_getscope(3C) 495
pthread_attr_getstack(3C) 496
pthread_attr_getstackaddr(3C) 498
pthread_attr_getstacksize(3C) 499
pthread_attr_init(3C) 500
pthread_barrierattr_destroy(3C) 502
pthread_barrierattr_getpshared(3C) 503
pthread_barrier_destroy(3C) 505
pthread_barrier_wait(3C) 507
pthread_cancel(3C) 509
pthread_cleanup_pop(3C) 510
pthread_cleanup_push(3C) 511
pthread_condattr_getclock(3C) 512
pthread_condattr_getpshared(3C) 513
pthread_condattr_init(3C) 515
pthread_cond_init(3C) 517
pthread_cond_signal(3C) 519
pthread_cond_wait(3C) 521
pthread_create(3C) 525
pthread_detach(3C) 529
pthread_equal(3C) 530
pthread_exit(3C) 531
pthread_getconcurrency(3C) 532
pthread_getschedparam(3C) 534
pthread_getspecific(3C) 536
pthread_join(3C) 537
pthread_key_create(3C) 539
pthread_key_delete(3C) 541
pthread_kill(3C) 542
pthread_mutexattr_getprioceiling(3C) 543
pthread_mutexattr_getprotocol(3C) 545

pthread_mutexattr_getpshared(3C) 548
 pthread_mutexattr_getrobust_np(3C) 550
 pthread_mutexattr_gettype(3C) 552
 pthread_mutexattr_init(3C) 554
 pthread_mutex_consistent_np(3C) 555
 pthread_mutex_getprioceiling(3C) 557
 pthread_mutex_init(3C) 559
 pthread_mutex_lock(3C) 561
 pthread_mutex_timedlock(3C) 564
 pthread_once(3C) 566
 pthread_rwlockattr_getpshared(3C) 567
 pthread_rwlockattr_init(3C) 568
 pthread_rwlock_init(3C) 569
 pthread_rwlock_rdlock(3C) 571
 pthread_rwlock_timedrdlock(3C) 573
 pthread_rwlock_timedwrlock(3C) 575
 pthread_rwlock_unlock(3C) 577
 pthread_rwlock_wrlock(3C) 578
 pthread_self(3C) 580
 pthread_setcancelstate(3C) 581
 pthread_setcanceltype(3C) 583
 pthread_setschedprio(3C) 585
 pthread_sigmask(3C) 586
 pthread_spin_destroy(3C) 591
 pthread_spin_lock(3C) 593
 pthread_spin_unlock(3C) 594
 pthread_testcancel(3C) 595
 ptrace(3C) 596
 ptsname(3C) 599
 putenv(3C) 600
 putpwent(3C) 601
 puts(3C) 602
 putspent(3C) 603
 putws(3C) 604
 qsort(3C) 605
 raise(3C) 607
 rand(3C) 608
 rand(3UCB) 609

random(3C) 610
rctlblk_set_value(3C) 613
rctl_walk(3C) 618
readdir(3C) 620
readdir(3UCB) 624
realpath(3C) 626
reboot(3C) 628
re_comp(3C) 629
regcmp(3C) 630
regcomp(3C) 632
remove(3C) 638
rewind(3C) 639
rewinddir(3C) 640
rwlock(3C) 641
scandir(3C) 644
scandir(3UCB) 646
scanf(3C) 647
schedctl_init(3C) 655
seekdir(3C) 657
select(3C) 658
semaphore(3C) 663
setbuf(3C) 667
setbuffer(3C) 669
setcat(3C) 670
setenv(3C) 671
setjmp(3C) 672
setjmp(3UCB) 675
setkey(3C) 678
setlabel(3C) 679
setlocale(3C) 680
sigblock(3UCB) 683
sigfpe(3C) 684
siginterrupt(3C) 686
siginterrupt(3UCB) 687
signal(3C) 688
signal(3UCB) 690
sigsetops(3C) 692
sigstack(3C) 694

sigstack(3UCB) 696
sigvec(3UCB) 697
sleep(3C) 702
sleep(3UCB) 703
ssignal(3C) 704
stack_getbounds(3C) 705
_stack_grow(3C) 706
stack_inbounds(3C) 707
stack_setbounds(3C) 708
stack_violation(3C) 709
stdio(3C) 711
str2sig(3C) 715
strcoll(3C) 716
strerror(3C) 717
strfmon(3C) 718
strftime(3C) 722
string(3C) 727
string_to_decimal(3C) 733
strptime(3C) 737
strsignal(3C) 743
strtod(3C) 744
strtoimax(3C) 749
strtol(3C) 750
strtoul(3C) 753
strtows(3C) 755
strxfrm(3C) 756
swab(3C) 758
sync_instruction_memory(3C) 759
syscall(3UCB) 760
sysconf(3C) 761
syslog(3C) 769
system(3C) 773
tcdrain(3C) 774
tcfLOW(3C) 775
tcflush(3C) 776
tcgetattr(3C) 777
tcgetpgrp(3C) 778
tcgetsid(3C) 779

tcsendbreak(3C) 780
tcsetattr(3C) 781
tcsetpgrp(3C) 783
td_init(3C_DB) 784
td_log(3C_DB) 785
td_sync_get_info(3C_DB) 786
td_ta_enable_stats(3C_DB) 790
td_ta_event_addr(3C_DB) 792
td_ta_get_nthreads(3C_DB) 796
td_ta_map_addr2sync(3C_DB) 797
td_ta_map_id2thr(3C_DB) 798
td_ta_new(3C_DB) 799
td_ta_setconcurrency(3C_DB) 801
td_ta_sync_iter(3C_DB) 802
td_thr_dbsuspend(3C_DB) 804
td_thr_getgregs(3C_DB) 805
td_thr_get_info(3C_DB) 807
td_thr_lockowner(3C_DB) 810
td_thr_setprio(3C_DB) 811
td_thr_setsigpending(3C_DB) 812
td_thr_sleepinfo(3C_DB) 813
td_thr_tsd(3C_DB) 814
td_thr_validate(3C_DB) 815
tell(3C) 816
telldir(3C) 817
termios(3C) 818
thr_create(3C) 819
thr_exit(3C) 825
thr_getconcurrency(3C) 827
thr_getprio(3C) 828
thr_join(3C) 829
thr_keycreate(3C) 831
thr_kill(3C) 834
thr_main(3C) 835
thr_min_stack(3C) 836
thr_self(3C) 837
thr_sigsetmask(3C) 838
thr_stksegment(3C) 842

thr_suspend(3C) 843
thr_yield(3C) 844
times(3UCB) 845
tmpfile(3C) 846
tmpnam(3C) 847
toascii(3C) 849
_tolower(3C) 850
tolower(3C) 851
_toupper(3C) 852
toupper(3C) 853
towctrans(3C) 854
tolower(3C) 855
toupper(3C) 856
truncate(3C) 857
tsearch(3C) 860
ttyname(3C) 863
ttyslot(3C) 865
ualarm(3C) 866
ucred_get(3C) 867
umem_alloc(3MALLOC) 870
umem_cache_create(3MALLOC) 875
umem_debug(3MALLOC) 883
ungetc(3C) 886
ungetwc(3C) 887
unlockpt(3C) 888
unsetenv(3C) 889
usleep(3C) 890
vfwprintf(3C) 891
vlfmt(3C) 892
vpfmt(3C) 894
vprintf(3C) 896
vsyslog(3C) 898
wait3(3C) 899
wait(3C) 902
wait(3UCB) 904
waitpid(3C) 908
walkcontext(3C) 910
watchmalloc(3MALLOC) 912

wcrtomb(3C) 915
wscoll(3C) 917
wcsftime(3C) 918
wcsrtombs(3C) 919
wcsstr(3C) 921
wcstod(3C) 922
wcstoimax(3C) 925
wcstol(3C) 926
wcstombs(3C) 929
wcstoul(3C) 930
wcstring(3C) 932
wcswidth(3C) 937
wcsxfrm(3C) 938
wctob(3C) 940
wctomb(3C) 941
wctrans(3C) 942
wctype(3C) 943
wcwidth(3C) 944
wmemchr(3C) 945
wmemcmp(3C) 946
wmemcpy(3C) 947
wmemmove(3C) 948
wmemset(3C) 949
wordexp(3C) 950
wsprintf(3C) 954
wscanf(3C) 955
wstring(3C) 956

Index 957

Preface

Both novice users and those familiar with the SunOS operating system can use online man pages to obtain information about the system and its features. 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.

Overview

The following contains a brief description of each man page section 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.
- 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 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 can 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	This section shows the syntax of commands or functions. When a command or file does not exist in the standard path, its full path name is shown. Options and arguments are alphabetized, with single letter arguments first, and options with arguments next, unless a different argument order is required.
	The following special characters are used in this section:
[]	Brackets. The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument must be specified.
. . .	Ellipses. Several values can be provided for the previous argument, or the previous argument can be specified multiple times, for example, "filename ...".
	Separator. Only one of the arguments separated by this character can be specified at a 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, and functions are described under USAGE.
IOCTL	This section appears on pages in Section 7 only. Only the device class that supplies appropriate parameters to the <code>ioctl(2)</code> system call is called <code>ioctl</code> and generates its own heading. <code>ioctl</code> calls for a specific device are listed alphabetically (on the man page for that specific device). <code>ioctl</code> calls are used for a particular class of devices all of which have an <code>io</code> ending, such as <code>mtio(7I)</code> .
OPTIONS	This section 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 <code>errno</code> indicating why they failed. This section lists alphabetically all error codes a function can generate and describes the

	<p>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.</p>
USAGE	<p>This section lists special rules, features, and commands that require in-depth explanations. The subsections listed here are used to explain built-in functionality:</p> <ul style="list-style-type: none"> Commands Modifiers Variables Expressions Input Grammar
EXAMPLES	<p>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 <code>example%</code>, or if the user must be superuser, <code>example#</code>. Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS, and USAGE sections.</p>
ENVIRONMENT VARIABLES	<p>This section lists any environment variables that the command or function affects, followed by a brief description of the effect.</p>
EXIT STATUS	<p>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.</p>
FILES	<p>This section lists all file names 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.</p>
ATTRIBUTES	<p>This section lists characteristics of commands, utilities, and device drivers by defining the attribute type and its corresponding value. See <code>attributes(5)</code> for more information.</p>
SEE ALSO	<p>This section lists references to other man pages, in-house documentation, and outside publications.</p>

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.

Basic Library Functions

a64l(3C)

NAME a64l, l64a – convert between long integer and base-64 ASCII string

SYNOPSIS

```
#include <stdlib.h>
long a64l(const char *s);
char *l64a(long l);
```

DESCRIPTION These functions maintain numbers stored in base-64 ASCII characters that define a notation by which long integers can be represented by up to six characters. Each character represents a “digit” in a radix-64 notation.

The characters used to represent “digits” are as follows:

Character	Digit
.	0
/	1
0-9	2-11
A-Z	12-37
a-z	38-63

The a64l () function takes a pointer to a null-terminated base-64 representation and returns a corresponding long value. If the string pointed to by s contains more than six characters, a64l () uses the first six.

The a64l () function scans the character string from left to right with the least significant digit on the left, decoding each character as a 6-bit radix-64 number.

The l64a () function takes a long argument and returns a pointer to the corresponding base-64 representation. If the argument is 0, l64a () returns a pointer to a null string.

The value returned by l64a () is a pointer into a static buffer, the contents of which are overwritten by each call. In the case of multithreaded applications, the return value is a pointer to thread specific data.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO attributes(5), standards(5)

NAME	abort – terminate the process abnormally						
SYNOPSIS	<pre>#include <stdlib.h> void abort(void);</pre>						
DESCRIPTION	<p>The <code>abort()</code> function causes abnormal process termination to occur, unless the signal <code>SIGABRT</code> is being caught and the signal handler does not return. The abnormal termination processing includes at least the effect of <code>fclose(3C)</code> on all open streams and message catalogue descriptors, and the default actions defined for <code>SIGABRT</code>. The <code>SIGABRT</code> signal is sent to the calling process as if by means of the <code>raise(3C)</code> function with the argument <code>SIGABRT</code>.</p> <p>The status made available to <code>wait(3C)</code> or <code>waitpid(3C)</code> by <code>abort</code> will be that of a process terminated by the <code>SIGABRT</code> signal. <code>abort</code> will override blocking or ignoring the <code>SIGABRT</code> signal.</p>						
RETURN VALUES	The <code>abort()</code> function does not return.						
ERRORS	No errors are defined.						
USAGE	Catching the signal is intended to provide the application writer with a portable means to abort processing, free from possible interference from any implementation-provided library functions. If <code>SIGABRT</code> is neither caught nor ignored, and the current directory is writable, a core dump may be produced.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	<code>exit(2)</code> , <code>getrlimit(2)</code> , <code>kill(2)</code> , <code>fclose(3C)</code> , <code>raise(3C)</code> , <code>signal(3C)</code> , <code>wait(3C)</code> , <code>waitpid(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

abs(3C)

NAME abs, labs, llabs – return absolute value of integer

SYNOPSIS

```
#include <stdlib.h>
int abs(int val);
long labs(long lval);
long long llabs(long long llval);
```

DESCRIPTION The `abs()` function returns the absolute value of its `int` operand.
The `labs()` function returns the absolute value of its `long` operand.
The `llabs()` function returns the absolute value of its `long long` operand.

USAGE In 2's-complement representation, the absolute value of the largest magnitude negative integral value is undefined.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `attributes(5)`, `standards(5)`

NAME	addsev – define additional severities				
SYNOPSIS	<pre>#include <pfmt.h> int addsev(int <i>int_val</i>, const char *<i>string</i>);</pre>				
DESCRIPTION	<p>The <code>addsev()</code> function defines additional severities for use in subsequent calls to <code>pfmt(3C)</code> or <code>lfmt(3C)</code>. It associates an integer value <code>int_val</code> in the range [5-255] with a character <code>string</code>, overwriting any previous string association between <code>int_val</code> and <code>string</code>.</p> <p>If <code>int_val</code> is OR-ed with the <code>flags</code> argument passed to subsequent calls to <code>pfmt()</code> or <code>lfmt()</code>, <code>string</code> will be used as severity. Passing a null <code>string</code> removes the severity.</p>				
RETURN VALUES	Upon successful completion, <code>addsev()</code> returns 0. Otherwise it returns -1.				
USAGE	Only the standard severities are automatically displayed for the locale in effect at runtime. An application must provide the means for displaying locale-specific versions of add-on severities. Add-on severities are only effective within the applications defining them.				
EXAMPLES	<p>EXAMPLE 1 Example of <code>addsev()</code> function.</p> <p>The following example</p> <pre>#define Panic 5 setlabel("APPL"); setcat("my_appl"); addsev(Panic, gettext(":26", "PANIC")); /* . . . */ lfmt(stderr, MM_SOFT MM_APPL PANIC, ":12:Cannot locate database\n");</pre> <p>will display the message to <code>stderr</code> and forward to the logging service</p> <pre>APPL: PANIC: Cannot locate database</pre>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<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>MT-safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-safe				
SEE ALSO	<code>gettext(3C)</code> , <code>lfmt(3C)</code> , <code>pfmt(3C)</code> , <code>attributes(5)</code>				

addseverity(3C)

NAME	addseverity – build a list of severity levels for an application for use with <code>fmtmsg</code>										
SYNOPSIS	<pre>#include <fmtmsg.h> int addseverity(int <i>severity</i>, const char *<i>string</i>);</pre>										
DESCRIPTION	<p>The <code>addseverity()</code> function builds a list of severity levels for an application to be used with the message formatting facility <code>fmtmsg()</code>. The <i>severity</i> argument is an integer value indicating the seriousness of the condition. The <i>string</i> argument is a pointer to a string describing the condition (string is not limited to a specific size).</p> <p>If <code>addseverity()</code> is called with an integer value that has not been previously defined, the function adds that new severity value and print string to the existing set of standard severity levels.</p> <p>If <code>addseverity()</code> is called with an integer value that has been previously defined, the function redefines that value with the new print string. Previously defined severity levels may be removed by supplying the null string. If <code>addseverity()</code> is called with a negative number or an integer value of 0, 1, 2, 3, or 4, the function fails and returns -1. The values 0-4 are reserved for the standard severity levels and cannot be modified. Identifiers for the standard levels of severity are:</p> <table><tr><td><code>MM_HALT</code></td><td>Indicates that the application has encountered a severe fault and is halting. Produces the print string <code>HALT</code>.</td></tr><tr><td><code>MM_ERROR</code></td><td>Indicates that the application has detected a fault. Produces the print string <code>ERROR</code>.</td></tr><tr><td><code>MM_WARNING</code></td><td>Indicates a condition that is out of the ordinary, that might be a problem, and should be watched. Produces the print string <code>WARNING</code>.</td></tr><tr><td><code>MM_INFO</code></td><td>Provides information about a condition that is not in error. Produces the print string <code>INFO</code>.</td></tr><tr><td><code>MM_NOSEV</code></td><td>Indicates that no severity level is supplied for the message.</td></tr></table> <p>Severity levels may also be defined at run time using the <code>SEV_LEVEL</code> environment variable (see <code>fmtmsg(3C)</code>).</p>	<code>MM_HALT</code>	Indicates that the application has encountered a severe fault and is halting. Produces the print string <code>HALT</code> .	<code>MM_ERROR</code>	Indicates that the application has detected a fault. Produces the print string <code>ERROR</code> .	<code>MM_WARNING</code>	Indicates a condition that is out of the ordinary, that might be a problem, and should be watched. Produces the print string <code>WARNING</code> .	<code>MM_INFO</code>	Provides information about a condition that is not in error. Produces the print string <code>INFO</code> .	<code>MM_NOSEV</code>	Indicates that no severity level is supplied for the message.
<code>MM_HALT</code>	Indicates that the application has encountered a severe fault and is halting. Produces the print string <code>HALT</code> .										
<code>MM_ERROR</code>	Indicates that the application has detected a fault. Produces the print string <code>ERROR</code> .										
<code>MM_WARNING</code>	Indicates a condition that is out of the ordinary, that might be a problem, and should be watched. Produces the print string <code>WARNING</code> .										
<code>MM_INFO</code>	Provides information about a condition that is not in error. Produces the print string <code>INFO</code> .										
<code>MM_NOSEV</code>	Indicates that no severity level is supplied for the message.										
RETURN VALUES	Upon successful completion, <code>addseverity()</code> returns <code>MM_OK</code> . Otherwise it returns <code>MM_NOTOK</code> .										
EXAMPLES	<p>EXAMPLE 1 Example of <code>addseverity()</code> function.</p> <p>When the function call</p> <pre>addseverity(7, "ALERT")</pre> <p>is followed by the call</p> <pre>fmtmsg(MM_PRINT, "UX:cat", 7, "invalid syntax", "refer to manual", "UX:cat:001")</pre>										

addseverity(3C)

EXAMPLE 1 Example of addseverity() function. (Continued)

the resulting output is

```
UX:cat: ALERT: invalid syntax
TO FIX: refer to manual   UX:cat:001
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO [fmtmsg\(1\)](#), [fmtmsg\(3C\)](#), [gettext\(3C\)](#), [printf\(3C\)](#), [attributes\(5\)](#)

assert(3C)

NAME	assert – verify program assertion						
SYNOPSIS	<pre>#include <assert.h> void assert (int <i>expression</i>);</pre>						
DESCRIPTION	<p>The <code>assert()</code> macro inserts diagnostics into applications. When executed, if <i>expression</i> is FALSE (zero), <code>assert()</code> prints the error message</p> <pre>Assertion failed: <i>expression</i>, file <i>xyz</i>, line <i>nnn</i></pre> <p>on the standard error output and aborts. In the error message, <i>xyz</i> is the name of the source file and <i>nnn</i> the source line number of the <code>assert()</code> statement. These are respectively the values of the preprocessor macros <code>__FILE__</code> and <code>__LINE__</code>.</p> <p>Since <code>assert()</code> is implemented as a macro, the <i>expression</i> may not contain any string literals.</p> <p>Compiling with the preprocessor option <code>-DNDEBUG</code> (see <code>cc(1B)</code>), or with the preprocessor control statement <code>#define NDEBUG</code> ahead of the <code>#include <assert.h></code> statement, will stop assertions from being compiled into the program.</p> <p>If the application is linked with <code>-lintl</code>, messages printed from this function are in the native language specified by the <code>LC_MESSAGES</code> locale category; see <code>setlocale(3C)</code>.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	<code>cc(1B)</code> , <code>abort(3C)</code> , <code>gettext(3C)</code> , <code>setlocale(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	atexit – register a function to run at process termination or object unloading						
SYNOPSIS	<pre>#include <stdlib.h> int atexit(void (*<i>func</i>)(void));</pre>						
DESCRIPTION	<p>The <code>atexit()</code> function registers the function pointed to by <i>func</i> to be called without arguments on normal termination of the program or when the object defining the function is unloaded.</p> <p>Normal termination occurs by either a call to the <code>exit(3C)</code> function or a return from <code>main()</code>. Object unloading occurs when a call to <code>dldclose(3C)</code> results in the object becoming unreferenced.</p> <p>The number of functions that may be registered with <code>atexit()</code> is limited only by available memory (refer to the <code>_SC_ATEXIT_MAX</code> argument of <code>sysconf(3C)</code>).</p> <p>After a successful call to any of the <code>exec(2)</code> functions, any functions previously registered by <code>atexit()</code> are no longer registered.</p> <p>On process exit, functions are called in the reverse order of their registration. On object unloading, any functions belonging to an unloadable object are called in the reverse order of their registration.</p>						
RETURN VALUES	Upon successful completion, the <code>atexit()</code> function returns 0. Otherwise, it returns a non-zero value.						
ERRORS	<p>The <code>atexit()</code> function may fail if:</p> <p><code>ENOMEM</code> Insufficient storage space is available.</p>						
USAGE	<p>The functions registered by a call to <code>atexit()</code> must return to ensure that all registered functions are called.</p> <p>There is no way for an application to tell how many functions have already been registered with <code>atexit()</code>.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	<code>exec(2)</code> , <code>dldclose(3C)</code> , <code>exit(3C)</code> , <code>sysconf(3C)</code> , <code>attributes(5)</code>						

atomic_ops(3C)

NAME atomic_ops, atomic_add_16, atomic_add_32, atomic_add_64, atomic_add_long, atomic_or_uint, atomic_or_32, atomic_and_uint, atomic_and_32, atomic_add_16_nv, atomic_add_32_nv, atomic_add_64_nv, atomic_add_long_nv – atomic operations

SYNOPSIS #include <atomic.h>

```
void atomic_add_16(uint16_t *target, int16_t delta);
void atomic_add_32(uint32_t *target, int32_t delta);
void atomic_add_long(ulong_t *target, long delta);
void atomic_add_64(uint64_t *target, uint64_t delta);
void atomic_or_uint(uint_t *target, uint_t bits);
void atomic_or_32(uint32_t *target, uint32_t bits);
void atomic_and_uint(uint_t *target, uint_t bits);
void atomic_and_32(uint32_t *target, uint32_t bits);
uint16_t atomic_add_16_nv(uint16_t *target, int16_t delta);
uint32_t atomic_add_32_nv(uint32_t *target, int32_t delta);
ulong_t atomic_add_long_nv(ulong_t *target, long delta);
uint64_t atomic_add_64_nv(uint64_t *target, int64_t delta);
```

DESCRIPTION The atomic_add_16(), atomic_add_32(), atomic_add_64(), and atomic_add_long() functions add *delta* to *target*.

The atomic_or_uint() and atomic_or_32() functions perform the bitwise-inclusive OR operation of *bits* with *target*.

The atomic_and_uint() and atomic_and_32() functions perform the bitwise AND operation of *bits* with *target*.

The atomic_add_16_nv(), atomic_add_32_nv(), atomic_add_long_nv(), and atomic_add_64_nv() functions behave as their counterparts atomic_add_16(), atomic_add_32(), atomic_add_64(), and atomic_add_long(), respectively, except that they also return the new value.

RETURN VALUES The atomic_add_16_nv(), atomic_add_32_nv(), atomic_add_long_nv(), and atomic_add_64_nv() functions return the new value of *target*.

ERRORS No errors are defined.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Stable

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO attributes(5)

NOTES The `atomic_add_16_nv()`, `atomic_add_32_nv()`, `atomic_add_long_nv()`, and `atomic_add_64_nv()` functions are substantially more expensive on some platforms than `atomic_add_16()`, `atomic_add_32()`, `atomic_add_64()`, and `atomic_add_long()`. Applications should avoid using the `*_nv()` functions unless there is a need to know the new value atomically, for example, when decrementing a reference count and checking whether it reached 0.

attropen(3C)

NAME attropen – open a file

SYNOPSIS

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

int attropen(const char *path, const char *attrpath, int oflag, /*
    mode_t mode */...);
```

DESCRIPTION

The `attropen()` function is similar to the `open(2)` function except that it takes a second path argument, `attrpath`, that identifies an extended attribute file associated with the first `path` argument. This function returns a file descriptor for the extended attribute rather than the file named by the initial argument.

The `O_XATTR` flag is set by default for `attropen()` and the `attrpath` argument is always interpreted as a reference to an extended attribute. Extended attributes must be referenced with a relative path; providing an absolute path results in a normal file reference.

RETURN VALUES Refer to `open(2)`.

ERRORS Refer to `open(2)`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving
MT-Level	Safe

SEE ALSO `open(2)`, `attributes(5)`, `fsattr(5)`

NAME basename – return the last element of a path name

SYNOPSIS #include <libgen.h>
char ***basename**(char **path*);

DESCRIPTION The `basename()` function takes the pathname pointed to by *path* and returns a pointer to the final component of the pathname, deleting any trailing '/' characters.

If the string consists entirely of the '/' character, `basename()` returns a pointer to the string "/".

If *path* is a null pointer or points to an empty string, `basename()` returns a pointer to the string ".".

RETURN VALUES The `basename()` function returns a pointer to the final component of *path*.

USAGE The `basename()` function may modify the string pointed to by *path*, and may return a pointer to static storage that may then be overwritten by a subsequent call to `basename()`.

When compiling multithreaded applications, the `_REENTRANT` flag must be defined on the compile line. This flag should only be used in multithreaded applications.

EXAMPLES **EXAMPLE 1** Examples for Input String and Output String

Input String	Output String
"/usr/lib"	"lib"
"/usr/"	"usr"
"/"	"/"

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `basename(1)`, `dirname(3C)`, `attributes(5)`, `standards(5)`

bsdmalloc(3MALLOC)

NAME	bsdmalloc – memory allocator				
SYNOPSIS	<pre>cc [flag ...] file ... -lbsdmalloc [library ...] char *malloc (size) ; unsigned size ; int free (ptr) ; char *ptr ; char *realloc (ptr , size) ; char *ptr ; unsigned size ;</pre>				
DESCRIPTION	<p>These routines provide a general-purpose memory allocation package. They maintain a table of free blocks for efficient allocation and coalescing of free storage. When there is no suitable space already free, the allocation routines call <code>sbrk(2)</code> to get more memory from the system. Each of the allocation routines returns a pointer to space suitably aligned for storage of any type of object. Each returns a null pointer if the request cannot be completed.</p> <p>The <code>malloc()</code> function returns a pointer to a block of at least <i>size</i> bytes, which is appropriately aligned.</p> <p>The <code>free()</code> function releases a previously allocated block. Its argument is a pointer to a block previously allocated by <code>malloc()</code> or <code>realloc()</code>.</p> <p>The <code>realloc()</code> function changes the size of the block pointed to by <i>ptr</i> to <i>size</i> bytes and returns a pointer to the (possibly moved) block. The contents will be unchanged up to the lesser of the new and old sizes. If the new size of the block requires movement of the block, the space for the previous instantiation of the block is freed. If the new size is larger, the contents of the newly allocated portion of the block are unspecified. If <i>ptr</i> is <code>NULL</code>, <code>realloc()</code> behaves like <code>malloc()</code> for the specified size. If <i>size</i> is 0 and <i>ptr</i> is not a null pointer, the space pointed to is freed.</p>				
RETURN VALUES	The <code>malloc()</code> and <code>realloc()</code> functions return a null pointer if there is not enough available memory. When <code>realloc()</code> returns <code>NULL</code> , the block pointed to by <i>ptr</i> is left intact.				
ERRORS	If <code>malloc()</code> or <code>realloc()</code> returns unsuccessfully, <code>errno</code> will be set to indicate the following: <table><tr><td>ENOMEM</td><td><i>size</i> bytes of memory cannot be allocated because it exceeds the physical limits of the system.</td></tr><tr><td>EAGAIN</td><td>There is not enough memory available at this point in time to allocate <i>size</i> bytes of memory; but the application could try again later.</td></tr></table>	ENOMEM	<i>size</i> bytes of memory cannot be allocated because it exceeds the physical limits of the system.	EAGAIN	There is not enough memory available at this point in time to allocate <i>size</i> bytes of memory; but the application could try again later.
ENOMEM	<i>size</i> bytes of memory cannot be allocated because it exceeds the physical limits of the system.				
EAGAIN	There is not enough memory available at this point in time to allocate <i>size</i> bytes of memory; but the application could try again later.				
SEE ALSO	<code>brk(2)</code> , <code>malloc(3C)</code> , <code>malloc(3MALLOC)</code> , <code>mapmalloc(3MALLOC)</code>				

WARNINGS	<p>Use of <code>libbsdmalloc</code> renders an application non-SCD compliant.</p> <p>The <code>libbsdmalloc</code> routines are incompatible with the memory allocation routines in the standard C-library (<code>libc</code>): <code>malloc(3C)</code>, <code>alloca(3C)</code>, <code>calloc(3C)</code>, <code>free(3C)</code>, <code>memalign(3C)</code>, <code>realloc(3C)</code>, and <code>valloc(3C)</code>.</p>
NOTES	<p>Using <code>realloc()</code> with a block freed before the most recent call to <code>malloc()</code> or <code>realloc()</code> results in an error.</p> <p>The <code>malloc()</code> and <code>realloc()</code> functions return a non-null pointer if <i>size</i> is 0. These pointers should not be dereferenced.</p> <p>Always cast the value returned by <code>malloc()</code> and <code>realloc()</code>.</p> <p>Comparative features of <code>bsdmalloc</code>, <code>malloc(3MALLOC)</code>, and <code>malloc(3C)</code>:</p> <ul style="list-style-type: none">■ The <code>bsdmalloc()</code> routines afford better performance but are space-inefficient.■ The <code>malloc(3MALLOC)</code> routines are space-efficient but have slower performance.■ The standard, fully SCD-compliant <code>malloc(3C)</code> routines are a trade-off between performance and space-efficiency. <p>The <code>free()</code> function does not set <code>errno</code>.</p>

bsd_signal(3C)

NAME	bsd_signal – simplified signal facilities				
SYNOPSIS	<pre>#include <signal.h> void (*bsd_signal(int sig, void (*func) (int))) (int);</pre>				
DESCRIPTION	<p>The <code>bsd_signal()</code> function provides a partially compatible interface for programs written to historical system interfaces (see USAGE below).</p> <p>The function call <code>bsd_signal(sig, func)</code> has an effect as if implemented as:</p> <pre>void (*bsd_signal(int sig, void (*func) (int))) (int) { struct sigaction act, oact; act.sa_handler = func; act.sa_flags = SA_RESTART; sigemptyset(&act.sa_mask); sigaddset(&act.sa_mask, sig); if (sigaction(sig, &act, &oact) == -1) return(SIG_ERR); return(oact.sa_handler); }</pre> <p>The handler function should be declared:</p> <pre>void handler(int sig);</pre> <p>where <i>sig</i> is the signal number. The behavior is undefined if <i>func</i> is a function that takes more than one argument, or an argument of a different type.</p>				
RETURN VALUES	Upon successful completion, <code>bsd_signal()</code> returns the previous action for <i>sig</i> . Otherwise, <code>SIG_ERR</code> is returned and <code>errno</code> is set to indicate the error.				
ERRORS	Refer to sigaction(2) .				
USAGE	This function is a direct replacement for the BSD signal(3UCB) function for simple applications that are installing a single-argument signal handler function. If a BSD signal handler function is being installed that expects more than one argument, the application has to be modified to use sigaction(2) . The <code>bsd_signal()</code> function differs from signal(3UCB) in that the <code>SA_RESTART</code> flag is set and the <code>SA_RESETHAND</code> will be clear when <code>bsd_signal()</code> is used. The state of these flags is not specified for signal(3UCB) .				
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>Interface Stability</td><td>Standard</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Standard				
SEE ALSO	sigaction(2) , sigaddset(3C) , sigemptyset(3C) , signal(3UCB) , attributes(5) , standards(5)				

NAME	bsearch – binary search a sorted table
SYNOPSIS	<pre>#include <stdlib.h> void *bsearch(const void *key, const void *base, size_t nel, size_t size, int (*compar)(const void *,const void *));</pre>
DESCRIPTION	<p>The <code>bsearch()</code> function is a binary search routine generalized from Knuth (6.2.1) Algorithm B. It returns a pointer into a table (an array) indicating where a datum may be found or a null pointer if the datum cannot be found. The table must be previously sorted in increasing order according to a comparison function pointed to by <code>compar</code>.</p> <p>The <code>key</code> argument points to a datum instance to be sought in the table. The <code>base</code> argument points to the element at the base of the table. The <code>nel</code> argument is the number of elements in the table. The <code>size</code> argument is the number of bytes in each element.</p> <p>The comparison function pointed to by <code>compar</code> is called with two arguments that point to the <code>key</code> object and to an array element, in that order. The function must return an integer less than, equal to, or greater than 0 if the <code>key</code> object is considered, respectively, to be less than, equal to, or greater than the array element.</p>
RETURN VALUES	The <code>bsearch()</code> function returns a pointer to a matching member of the array, or a null pointer if no match is found. If two or more members compare equal, which member is returned is unspecified.
USAGE	<p>The pointers to the key and the element at the base of the table should be of type pointer-to-element.</p> <p>The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.</p> <p>If the number of elements in the table is less than the size reserved for the table, <code>nel</code> should be the lower number.</p> <p>The <code>bsearch()</code> function safely allows concurrent access by multiple threads to disjoint data, such as overlapping subtrees or tables.</p>
EXAMPLES	<p>EXAMPLE 1 Examples for searching a table containing pointers to nodes.</p> <p>The example below searches a table containing pointers to nodes consisting of a string and its length. The table is ordered alphabetically on the string in the node pointed to by each entry.</p> <p>This program reads in strings and either finds the corresponding node and prints out the string and its length, or prints an error message.</p> <pre>#include <stdio.h> #include <stdlib.h> #include <string.h> struct node { /* these are stored in the table */ char *string;</pre>

bsearch(3C)

EXAMPLE 1 Examples for searching a table containing pointers to nodes. (Continued)

```

    int length;
};
static struct node table[] = { /* table to be searched */
    { "asparagus", 10 },
    { "beans", 6 },
    { "tomato", 7 },
    { "watermelon", 11 },
};

main()
{
    struct node *node_ptr, node;
    /* routine to compare 2 nodes */
    static int node_compare(const void *, const void *);
    char str_space[20]; /* space to read string into */

    node.string = str_space;
    while (scanf("%20s", node.string) != EOF) {
        node_ptr = bsearch( &node,
            table, sizeof(table)/sizeof(struct node),
            sizeof(struct node), node_compare);
        if (node_ptr != NULL) {
            (void) printf("string = %20s, length = %d\n",
                node_ptr->string, node_ptr->length);
        } else {
            (void) printf("not found: %20s\n", node.string);
        }
    }
    return(0);
}

/* routine to compare two nodes based on an */
/* alphabetical ordering of the string field */
static int
node_compare(const void *node1, const void *node2) {
    return (strcmp(
        ((const struct node *)node1)->string,
        ((const struct node *)node2)->string));
}

```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [bsearch\(3C\)](#), [lsearch\(3C\)](#), [qsort\(3C\)](#), [tsearch\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	bstring, bcopy, bcmp, bzero – memory operations						
SYNOPSIS	<pre>#include <strings.h> void bcopy(const void *s1, void *s2, size_t n); int bcmp(const void *s1, const void *s2, size_t n); void bzero(void *s, size_t n);</pre>						
DESCRIPTION	<p>The <code>bcopy()</code>, <code>bcmp()</code>, and <code>bzero()</code> functions operate as efficiently as possible on memory areas (arrays of bytes bounded by a count, not terminated by a null character). They do not check for the overflow of any receiving memory area. These functions are similar to the <code>memcpy()</code>, <code>memcmp()</code>, and <code>memset()</code> functions described on the memory(3C) manual page.</p> <p>The <code>bcopy()</code> function copies n bytes from memory area $s1$ to $s2$. Copying between objects that overlap will take place correctly.</p> <p>The <code>bcmp()</code> function compares the first n bytes of its arguments, returning 0 if they are identical and 1 otherwise. The <code>bcmp()</code> function always returns 0 when n is 0.</p> <p>The <code>bzero()</code> function sets the first n bytes in memory area s to 0.</p>						
WARNINGS	The <code>bcopy()</code> function takes parameters backwards from <code>memcmp()</code> . See memory(3C) .						
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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	memory(3C) , attributes(5) , standards(5)						

btowc(3C)

NAME	btowc – single-byte to wide-character conversion						
SYNOPSIS	<pre>#include <stdio.h> #include <wchar.h> wint_t btowc (int c);</pre>						
DESCRIPTION	<p>The <code>btowc()</code> function determines whether <code>c</code> constitutes a valid (one-byte) character in the initial shift state.</p> <p>The behavior of this function is affected by the <code>LC_CTYPE</code> category of the current locale. See <code>environ(5)</code>.</p>						
RETURN VALUES	The <code>btowc()</code> function returns <code>WEOF</code> if <code>c</code> has the value <code>EOF</code> or if <code>(unsigned char)c</code> does not constitute a valid (one-byte) character in the initial shift state. Otherwise, it returns the wide-character representation of that character.						
ERRORS	No errors are defined.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe with exceptions</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe with exceptions						
SEE ALSO	<code>setlocale(3C)</code> , <code>wctob(3C)</code> , <code>attributes(5)</code> , <code>environ(5)</code> , <code>standards(5)</code>						
NOTES	The <code>btowc()</code> function can be used safely in multithreaded applications, as long as <code>setlocale(3C)</code> is not being called to change the locale.						

NAME	catgets – read a program message						
SYNOPSIS	<pre>#include <nl_types.h> char *catgets(nl_catd catd, int set_num, int msg_num, const char *s);</pre>						
DESCRIPTION	The <code>catgets()</code> function attempts to read message <code>msg_num</code> , in set <code>set_num</code> , from the message catalog identified by <code>catd</code> . The <code>catd</code> argument is a catalog descriptor returned from an earlier call to <code>catopen()</code> . The <code>s</code> argument points to a default message string which will be returned by <code>catgets()</code> if the identified message catalog is not currently available.						
RETURN VALUES	If the identified message is retrieved successfully, <code>catgets()</code> returns a pointer to an internal buffer area containing the null terminated message string. If the call is unsuccessful for any reason, <code>catgets()</code> returns a pointer to <code>s</code> and <code>errno</code> may be set to indicate the error.						
ERRORS	<p>The <code>catgets()</code> function may fail if:</p> <p>EBADF The <code>catd</code> argument is not a valid message catalogue descriptor open for reading.</p> <p>EBADMSG The number of <code>%n</code> specifiers that appear in the message string specified by <code>s</code> does not match the number of <code>%n</code> specifiers that appear in the message identified by <code>set_id</code> and <code>msg_id</code> in the specified message catalog.</p> <p>EINTR The read operation was terminated due to the receipt of a signal, and no data was transferred.</p> <p>EINVAL The message catalog identified by <code>catd</code> is corrupted.</p> <p>ENOMSG The message identified by <code>set_id</code> and <code>msg_id</code> is not in the message catalog.</p>						
USAGE	The <code>catgets()</code> function can be used safely in multithreaded applications as long as <code>setlocale(3C)</code> is not being called to change the locale.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<p><code>gencat(1)</code>, <code>catclose(3C)</code>, <code>catopen(3C)</code>, <code>gettext(3C)</code>, <code>setlocale(3C)</code>, <code>attributes(5)</code>, <code>standards(5)</code></p> <p><i>International Language Environments Guide</i></p>						

catopen(3C)

NAME	catopen, catclose – open/close a message catalog												
SYNOPSIS	<pre>#include <nl_types.h> nl_catd catopen(const char *name, int oflag); int catclose(nl_catd catd);</pre>												
DESCRIPTION	<p>The <code>catopen()</code> function opens a message catalog and returns a message catalog descriptor. <i>name</i> specifies the name of the message catalog to be opened. If <i>name</i> contains a “/”, then <i>name</i> specifies a complete pathname for the message catalog; otherwise, the environment variable <code>NLSPATH</code> is used and <code>/usr/lib/locale/locale/LC_MESSAGES</code> must exist. If <code>NLSPATH</code> does not exist in the environment, or if a message catalog cannot be opened in any of the paths specified by <code>NLSPATH</code>, then the default path <code>/usr/lib/locale/locale/LC_MESSAGES</code> is used. In the “C” locale, <code>catopen()</code> will always succeed without checking the default search path.</p> <p>The names of message catalogs and their location in the filesystem can vary from one system to another. Individual applications can choose to name or locate message catalogs according to their own special needs. A mechanism is therefore required to specify where the catalog resides.</p> <p>The <code>NLSPATH</code> variable provides both the location of message catalogs, in the form of a search path, and the naming conventions associated with message catalog files. For example:</p> <pre>NLSPATH=/nlslib/%L/%N.cat:/nlslib/%N/%L</pre> <p>The metacharacter <code>%</code> introduces a substitution field, where <code>%L</code> substitutes the current setting of either the <code>LANG</code> environment variable, if the value of <i>oflag</i> is 0, or the <code>LC_MESSAGES</code> category, if the value of <i>oflag</i> is <code>NL_CAT_LOCALE</code>, and <code>%N</code> substitutes the value of the <i>name</i> parameter passed to <code>catopen()</code>. Thus, in the above example, <code>catopen()</code> will search in <code>/nlslib/\$LANG/name.cat</code>, if <i>oflag</i> is 0, or in <code>/nlslib/{LC_MESSAGES}/name.cat</code>, if <i>oflag</i> is <code>NL_CAT_LOCALE</code>.</p> <p>The <code>NLSPATH</code> variable will normally be set up on a system wide basis (in <code>/etc/profile</code>) and thus makes the location and naming conventions associated with message catalogs transparent to both programs and users.</p> <p>The full set of metacharacters is:</p> <table><tr><td><code>%N</code></td><td>The value of the name parameter passed to <code>catopen()</code>.</td></tr><tr><td><code>%L</code></td><td>The value of <code>LANG</code> or <code>LC_MESSAGES</code>.</td></tr><tr><td><code>%l</code></td><td>The value of the <i>language</i> element of <code>LANG</code> or <code>LC_MESSAGES</code>.</td></tr><tr><td><code>%t</code></td><td>The value of the <i>territory</i> element of <code>LANG</code> or <code>LC_MESSAGES</code>.</td></tr><tr><td><code>%c</code></td><td>The value of the <i>codeset</i> element of <code>LANG</code> or <code>LC_MESSAGES</code>.</td></tr><tr><td><code>%%</code></td><td>A single <code>%</code>.</td></tr></table>	<code>%N</code>	The value of the name parameter passed to <code>catopen()</code> .	<code>%L</code>	The value of <code>LANG</code> or <code>LC_MESSAGES</code> .	<code>%l</code>	The value of the <i>language</i> element of <code>LANG</code> or <code>LC_MESSAGES</code> .	<code>%t</code>	The value of the <i>territory</i> element of <code>LANG</code> or <code>LC_MESSAGES</code> .	<code>%c</code>	The value of the <i>codeset</i> element of <code>LANG</code> or <code>LC_MESSAGES</code> .	<code>%%</code>	A single <code>%</code> .
<code>%N</code>	The value of the name parameter passed to <code>catopen()</code> .												
<code>%L</code>	The value of <code>LANG</code> or <code>LC_MESSAGES</code> .												
<code>%l</code>	The value of the <i>language</i> element of <code>LANG</code> or <code>LC_MESSAGES</code> .												
<code>%t</code>	The value of the <i>territory</i> element of <code>LANG</code> or <code>LC_MESSAGES</code> .												
<code>%c</code>	The value of the <i>codeset</i> element of <code>LANG</code> or <code>LC_MESSAGES</code> .												
<code>%%</code>	A single <code>%</code> .												

The LANG environment variable provides the ability to specify the user's requirements for native languages, local customs and character set, as an ASCII string in the form

```
LANG=language[_territory[.codeset]]
```

A user who speaks German as it is spoken in Austria and has a terminal which operates in ISO 8859/1 codeset, would want the setting of the LANG variable to be

```
LANG=De_A.88591
```

With this setting it should be possible for that user to find any relevant catalogs should they exist.

Should the LANG variable not be set, the value of LC_MESSAGES as returned by `setlocale()` is used. If this is NULL, the default path as defined in `<nl_types.h>` is used.

A message catalogue descriptor remains valid in a process until that process closes it, or a successful call to one of the `exec` functions. A change in the setting of the LC_MESSAGES category may invalidate existing open catalogues.

If a file descriptor is used to implement message catalogue descriptors, the FD_CLOEXEC flag will be set; see `<fcntl.h>`.

If the value of *oflag* argument is 0, the LANG environment variable is used to locate the catalogue without regard to the LC_MESSAGES category. If the *oflag* argument is NL_CAT_LOCALE, the LC_MESSAGES category is used to locate the message catalogue.

The `catclose()` function closes the message catalog identified by *catd*. If a file descriptor is used to implement the type `nl_catd`, that file descriptor will be closed.

RETURN VALUES Upon successful completion, `catopen()` returns a message catalog descriptor for use on subsequent calls to `catgets()` and `catclose()`. Otherwise it returns `(nl_catd) -1`.

Upon successful completion, `catclose()` returns 0. Otherwise it returns `-1` and sets `errno` to indicate the error.

ERRORS The `catopen()` function may fail if:

EACCES	Search permission is denied for the component of the path prefix of the message catalogue or read permission is denied for the message catalogue.
EMFILE	There are OPEN_MAX file descriptors currently open in the calling process.
ENAMETOOLONG	The length of the pathname of the message catalogue 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.

catopen(3C)

ENFILE	Too many files are currently open in the system.
ENOENT	The message catalogue does not exist or the <i>name</i> argument points to an empty string.
ENOMEM	Insufficient storage space is available.
ENOTDIR	A component of the path prefix of the message catalogue is not a directory.

The `catclose()` function may fail if:

EBADF	The catalogue descriptor is not valid.
EINTR	The <code>catclose()</code> function was interrupted by a signal.

USAGE The `catopen()` and `catclose()` functions can be used safely in multithreaded applications, as long as `setlocale(3C)` is not being called to change the locale.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `gencat(1)`, `catgets(3C)`, `gettext(3C)`, `nl_types.h(3HEAD)`, `setlocale(3C)`, `attributes(5)`, `environ(5)`

NAME	cfgetispeed, cfgetospeed – get input and output baud rate
SYNOPSIS	<pre>#include <termios.h> speed_t cfgetispeed(const struct termios *termios_p); speed_t cfgetospeed(const struct termios *termios_p);</pre>
DESCRIPTION	<p>The <code>cfgetispeed()</code> function extracts the input baud rate from the <code>termios</code> structure to which the <code>termios_p</code> argument points.</p> <p>The <code>cfgetospeed()</code> function extracts the output baud rate from the <code>termios</code> structure to which the <code>termios_p</code> argument points.</p> <p>These functions returns exactly the value in the <code>termios</code> data structure, without interpretation.</p>
RETURN VALUES	<p>Upon successful completion, <code>cfgetispeed()</code> returns a value of type <code>speed_t</code> representing the input baud rate.</p> <p>Upon successful completion, <code>cfgetospeed()</code> returns a value of type <code>speed_t</code> representing the output baud rate.</p>
ERRORS	No errors are defined.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe, and Async-Signal-Safe

SEE ALSO [cfgetospeed\(3C\)](#), [tcgetattr\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#), [termio\(7I\)](#)

cfsetispeed(3C)

NAME	cfsetispeed, cfsetospeed – set input and output baud rate						
SYNOPSIS	<pre>#include <termios.h> int cfsetispeed(struct termios *termios_p, speed_t speed); int cfsetospeed(struct termios *termios_p, speed_t speed);</pre>						
DESCRIPTION	<p>The <code>cfsetispeed()</code> function sets the input baud rate stored in the structure pointed to by <code>termios_p</code> to <code>speed</code>.</p> <p>The <code>cfsetospeed()</code> function sets the output baud rate stored in the structure pointed to by <code>termios_p</code> to <code>speed</code>.</p> <p>There is no effect on the baud rates set in the hardware until a subsequent successful call to <code>tcsetattr(3C)</code> on the same <code>termios</code> structure.</p>						
RETURN VALUES	Upon successful completion, <code>cfsetispeed()</code> and <code>cfsetospeed()</code> return 0. Otherwise -1 is returned, and <code>errno</code> may be set to indicate the error.						
ERRORS	The <code>cfsetispeed()</code> and <code>cfsetospeed()</code> functions may fail if: EINVAL The <code>speed</code> value is not a valid baud rate. EINVAL The value of <code>speed</code> is outside the range of possible speed values as specified in <code><termios.h></code> .						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe, and Async-Signal-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe, and Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe, and Async-Signal-Safe						
SEE ALSO	<code>cfgetispeed(3C)</code> , <code>tcsetattr(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code> , <code>termio(7I)</code>						

NAME	clock – report CPU time used						
SYNOPSIS	<pre>#include <time.h> clock_t clock(void);</pre>						
DESCRIPTION	The <code>clock()</code> function returns the amount of CPU time (in microseconds) used since the first call to <code>clock()</code> in the calling process. The time reported is the sum of the user and system times of the calling process and its terminated child processes for which it has executed the <code>wait(3C)</code> function, the <code>pclose(3C)</code> function, or the <code>system(3C)</code> function.						
RETURN VALUES	Dividing the value returned by <code>clock()</code> by the constant <code>CLOCKS_PER_SEC</code> , defined in the <code><time.h></code> header, will give the time in seconds. If the process time used is not available or cannot be represented, <code>clock</code> returns the value <code>(clock_t) -1</code> .						
USAGE	The value returned by <code>clock()</code> is defined in microseconds for compatibility with systems that have CPU clocks with much higher resolution. Because of this, the value returned will wrap around after accumulating only 2147 seconds of CPU time (about 36 minutes).						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>times(2)</code> , <code>popen(3C)</code> , <code>system(3C)</code> , <code>wait(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

closedir(3C)

NAME	closedir – close a directory stream						
SYNOPSIS	<pre>#include <sys/types.h> #include <dirent.h> int closedir(DIR *dirp);</pre>						
DESCRIPTION	The <code>closedir()</code> function closes the directory stream referred to by the argument <i>dirp</i> . Upon return, the value of <i>dirp</i> may no longer point to an accessible object of the type DIR. If a file descriptor is used to implement type DIR, that file descriptor will be closed.						
RETURN VALUES	Upon successful completion, <code>closedir()</code> returns 0. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.						
ERRORS	The <code>closedir()</code> function may fail if: EBADF The <i>dirp</i> argument does not refer to an open directory stream. EINTR The <code>closedir()</code> function was interrupted by a signal.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	<code>opendir(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	closefrom, fdwalk – close or iterate over open file descriptors
SYNOPSIS	<pre>#include <stdlib.h> void closefrom(int <i>lowfd</i>) ; int fdwalk(int (<i>*func</i>)(void *, int), void *<i>cd</i>) ;</pre>
DESCRIPTION	<p>The <code>closefrom()</code> function calls <code>close(2)</code> on all open file descriptors greater than or equal to <i>lowfd</i>.</p> <p>The effect of <code>closefrom(<i>lowfd</i>)</code> is the same as the code</p> <pre>#include <sys/resource.h> struct rlimit rl; int i; getrlimit(RLIMIT_NOFILE, &rl); for (i = lowfd; i < rl.rlim_max; i++) (void) close(i);</pre> <p>except that <code>close()</code> is called only on file descriptors that are actually open, not on every possible file descriptor greater than or equal to <i>lowfd</i>, and <code>close()</code> is also called on any open file descriptors greater than or equal to <code>rl.rlim_max</code> (and <i>lowfd</i>), should any exist.</p> <p>The <code>fdwalk()</code> function first makes a list of all currently open file descriptors. Then for each file descriptor in the list, it calls the user-defined function, <i>func</i>(<i>cd</i>, <i>fd</i>), passing it the pointer to the callback data, <i>cd</i>, and the value of the file descriptor from the list, <i>fd</i>. The list is processed in file descriptor value order, lowest numeric value first.</p> <p>If <i>func</i>() returns a non-zero value, the iteration over the list is terminated and <code>fdwalk()</code> returns the non-zero value returned by <i>func</i>(). Otherwise, <code>fdwalk()</code> returns 0 after having called <i>func</i>() for every file descriptor in the list.</p> <p>The <code>fdwalk()</code> function can be used for fine-grained control over the closing of file descriptors. For example, the <code>closefrom()</code> function can be implemented as:</p> <pre>static int close_func(void *lowfdp, int fd) { if (fd >= *(int *)lowfdp) (void) close(fd); return (0); } void closefrom(int lowfd) { (void) fdwalk(close_func, &lowfd); }</pre> <p>The <code>fdwalk()</code> function can then be used to count the number of open files in the process.</p>

closefrom(3C)

RETURN VALUES No return value is defined for `closefrom()`. If `close()` fails for any of the open file descriptors, the error is ignored and the file descriptors whose `close()` operation failed might remain open on return from `closefrom()`.

The `fdwalk()` function returns the return value of the last call to the callback function `func()`, or 0 if `func()` is never called (no open files).

ERRORS No errors are defined. The `closefrom()` and `fdwalk()` functions do not set `errno` but `errno` can be set by `close()` or by another function called by the callback function, `func()`.

FILES `/proc/self/fd` directory (list of open files)

USAGE The act of closing all open file descriptors should be performed only as the first action of a daemon process. Closing file descriptors that are in use elsewhere in the current process normally leads to disastrous results.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

SEE ALSO `close(2)`, `getrlimit(2)`, `proc(4)`, `attributes(5)`

NAME	cond_init, cond_wait, cond_timedwait, cond_reltimedwait, cond_signal, cond_broadcast, cond_destroy – condition variables
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <thread.h> #include <synch.h> int cond_init(cond_t *cvp, int type, void *arg); int cond_wait(cond_t *cvp, mutex_t *mp); int cond_timedwait(cond_t *cvp, mutex_t *mp, timestruc_t *abstime); int cond_reltimedwait(cond_t *cvp, mutex_t *mp, timestruc_t *reltime); int cond_signal(cond_t *cvp); int cond_broadcast(cond_t *cvp); int cond_destroy(cond_t *cvp);</pre>
DESCRIPTION	
Initialize	<p>Condition variables and mutexes should be global. Condition variables that are allocated in writable memory can synchronize threads among processes if they are shared by the cooperating processes (see <code>mmap(2)</code>) and are initialized for this purpose.</p> <p>The scope of a condition variable is either intra-process or inter-process. This is dependent upon whether the argument is passed implicitly or explicitly to the initialization of that condition variable. A condition variable does not need to be explicitly initialized. A condition variable is initialized with all zeros, by default, and its scope is set to within the calling process. For inter-process synchronization, a condition variable must be initialized once, and only once, before use.</p> <p>A condition variable must not be simultaneously initialized by multiple threads or re-initialized while in use by other threads.</p> <p>Attributes of condition variables can be set to the default or customized at initialization.</p> <p>The <code>cond_init()</code> function initializes the condition variable pointed to by <code>cvp</code>. A condition variable can have several different types of behavior, specified by <code>type</code>. No current type uses <code>arg</code> although a future type may specify additional behavior parameters with <code>arg</code>. The <code>type</code> argument can take one of the following values:</p> <p><code>USYNC_THREAD</code> The condition variable can synchronize threads only in this process. This is the default.</p> <p><code>USYNC_PROCESS</code> The condition variable can synchronize threads in this process and other processes. Only one process should initialize the condition variable. The object initialized with this attribute must be allocated in memory shared between processes, either in System V shared memory (see <code>shmop(2)</code>) or in memory mapped to a file (see</p>

cond_init(3C)

mmap(2)). It is illegal to initialize the object this way and to not allocate it in such shared memory.

Initializing condition variables can also be accomplished by allocating in zeroed memory, in which case, a *type* of `USYNC_THREAD` is assumed.

If default condition variable attributes are used, statically allocated condition variables can be initialized by the macro `DEFAULTCV`.

Default condition variable initialization (intra-process):

```
cond_t cvp;  
  
cond_init(&cvp, NULL, NULL); /*initialize condition variable with default*/
```

or

```
cond_init(&cvp, USYNC_THREAD, NULL);
```

or

```
cond_t cond = DEFAULTCV;
```

Customized condition variable initialization (inter-process):

```
cond_init(&cvp, USYNC_PROCESS, NULL); /* initialize cv with  
inter-process scope */
```

Condition Wait

The condition wait interface allows a thread to wait for a condition and atomically release the associated mutex that it needs to hold to check the condition. The thread waits for another thread to make the condition true and that thread's resulting call to signal and wakeup the waiting thread.

The `cond_wait()` function atomically releases the mutex pointed to by *mp* and causes the calling thread to block on the condition variable pointed to by *cvp*. The blocked thread may be awakened by `cond_signal()`, `cond_broadcast()`, or when interrupted by delivery of a UNIX signal or a `fork()`.

The `cond_wait()`, `cond_timedwait()`, and `cond_reltimedwait()` functions always return with the mutex locked and owned by the calling thread even when returning an error, except when the mutex is of `USYNC_PROCESS_ROBUST` type and has been left irrecoverable by the mutex's last owner. The `cond_wait()`, `cond_timedwait()`, and `cond_reltimedwait()` functions return the appropriate error value if they fail to internally reacquire the mutex.

Condition Signaling

A condition signal allows a thread to unblock a single thread waiting on the condition variable, whereas a condition broadcast allows a thread to unblock all threads waiting on the condition variable.

The `cond_signal()` function unblocks one thread that is blocked on the condition variable pointed to by *cvp*.

The `cond_broadcast()` function unblocks all threads that are blocked on the condition variable pointed to by *cvp*.

If no threads are blocked on the condition variable, then `cond_signal()` and `cond_broadcast()` have no effect.

Both functions should be called under the protection of the same mutex that is used with the condition variable being signaled. Otherwise, the condition variable may be signaled between the test of the associated condition and blocking in `cond_wait()`. This can cause an infinite wait.

Destroy The condition destroy functions destroy any state, but not the space, associated with the condition variable.

The `cond_destroy()` function destroys any state associated with the condition variable pointed to by *cvp*. The space for storing the condition variable is not freed.

RETURN VALUES Upon successful completion, these functions return 0. Otherwise, a non-zero value is returned to indicate the error.

ERRORS The `cond_timedwait()` and `cond_reltimedwait()` functions will fail if:

ETIME The time specified by *abstime* or *reltime* has passed.

These functions may fail if:

EFAULT The *cond*, *attr*, *cvp*, *arg*, *abstime*, or *mutex* argument points to an illegal address.

EINVAL Invalid argument. For `cond_init()`, *type* is not a recognized type. For `cond_timedwait()`, the number of nanoseconds is greater than or equal to 1,000,000,000.

ELOCKUNMAPPED The last owner of this mutex unmapped the mutex while holding the mutex. This mutex is now owned by the caller. See the description of the `mutex_lock()` function on the [mutex_init\(3C\)](#) manual page.

ENOTRECOVERABLE The mutex pointed to by *mp* is protecting the state that has been left irrecoverable when the mutex's last owner was not able to clean up the state. The mutex has not been acquired. See the description of the `mutex_lock()` function on the [mutex_init\(3C\)](#) manual page.

EOWNERDEAD The last owner of the mutex pointed to by *mp* died while holding the mutex. The mutex has been acquired. See the description of the `mutex_lock()` function on the [mutex_init\(3C\)](#) manual page.

The `cond_wait()` function may fail if:

cond_init(3C)

EINTR

Interrupted. The calling thread was awakened by the delivery of a UNIX signal.

EXAMPLES

EXAMPLE 1 Use `cond_wait()` in a loop to test some condition.

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

```
(void) mutex_lock(mp);
while (cond == FALSE) {
    (void) cond_wait(cvp, mp);
}
(void) mutex_unlock(mp);
```

EXAMPLE 2 Use `cond_timedwait()` in a loop to test some condition.

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

```
timestruc_t to;
...
(void) mutex_lock(mp);
to.tv_sec = time(NULL) + TIMEOUT;
to.tv_nsec = 0;
while (cond == FALSE) {
    err = cond_timedwait(cvp, mp, &to);
    if (err == ETIME) {
        /* timeout, do something */
        break;
    }
}
(void) mutex_unlock(mp);
```

EXAMPLE 3 Use `cond_reltimedwait()` in a loop to test some condition.

The `cond_reltimedwait()` function is normally used in a loop testing in some condition. It uses a relative timeout value as follows:

```
timestruc_t to;
...
(void) mutex_lock(mp);
while (cond == FALSE) {
    to.tv_sec = TIMEOUT;
    to.tv_nsec = 0;
    err = cond_reltimedwait(cvp, mp, &to);
    if (err == ETIME) {
        /* timeout, do something */
        break;
    }
}
(void) mutex_unlock(mp);
```


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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `fork(2)`, `mmap(2)`, `setitimer(2)`, `shmop(2)`, `mutex_init(3C)`, `signal(3C)`, `attributes(5)`, `condition(5)`, `mutex(5)`, `standards(5)`

NOTES The only policy currently supported is `SCHED_OTHER`. In Solaris, under the `SCHED_OTHER` policy, there is no established order in which threads are unblocked.

If more than one thread is blocked on a condition variable, the order in which threads are unblocked is determined by the scheduling policy. When each thread, unblocked as a result of a `cond_signal()` or `cond_broadcast()`, returns from its call to `cond_wait()` or `cond_timedwait()`, the thread owns the mutex with which it called `cond_wait()`, `cond_timedwait()`, or `cond_reltimedwait()`. The thread(s) that are unblocked compete for the mutex according to the scheduling policy and as if each had called `mutex_lock(3C)`.

When `cond_wait()` returns the value of the condition is indeterminate and must be reevaluated.

The `cond_timedwait()` and `cond_reltimedwait()` functions are similar to `cond_wait()`, except that the calling thread will not wait for the condition to become true past the absolute time specified by *abstime* or the relative time specified by *reltime*. Note that `cond_timedwait()` or `cond_reltimedwait()` might continue to block as it tries to reacquire the mutex pointed to by *mp*, which may be locked by another thread. If either `cond_timedwait()` or `cond_reltimedwait()` returns because of a timeout, it returns the error value `ETIME`.

confstr(3C)

NAME	confstr – get configurable variables
SYNOPSIS	<pre>#include <unistd.h> size_t confstr(int <i>name</i>, char *<i>buf</i>, size_t <i>len</i>);</pre>
DESCRIPTION	<p>The <code>confstr()</code> function provides a method for applications to get configuration-defined string values. Its use and purpose are similar to the sysconf(3C) function, but it is used where string values rather than numeric values are returned.</p> <p>The <i>name</i> argument represents the system variable to be queried.</p> <p>If <i>len</i> is not 0, and if <i>name</i> has a configuration-defined value, <code>confstr()</code> copies that value into the <i>len</i>-byte buffer pointed to by <i>buf</i>. If the string to be returned is longer than <i>len</i> bytes, including the terminating null, then <code>confstr()</code> truncates the string to <i>len</i>-1 bytes and null-terminates the result. The application can detect that the string was truncated by comparing the value returned by <code>confstr()</code> with <i>len</i>.</p> <p>If <i>len</i> is 0, <code>confstr()</code> still returns the integer value as defined below, but does not return the string.</p> <p>The <code>confstr()</code> function supports the following values for <i>name</i>, defined in <code><unistd.h></code>, for both SPARC and x86:</p> <ul style="list-style-type: none"><code>_CS_LFS64_CFLAGS</code> If <code>_LFS64_LARGEFILE</code> is defined in <code><unistd.h></code>, this value is the set of initial options to be given to the <code>cc</code> and <code>c89</code> utilities to build an application using the Large File Summit transitional compilation environment (see <code>lfcompile64(5)</code>).<code>_CS_LFS64_LDFLAGS</code> If <code>_LFS64_LARGEFILE</code> is defined in <code><unistd.h></code>, this value is the set of final options to be given to the <code>cc</code> and <code>c89</code> utilities to build an application using the Large File Summit transitional compilation environment (see <code>lfcompile64(5)</code>).<code>_CS_LFS64_LIBS</code> If <code>_LFS64_LARGEFILE</code> is defined in <code><unistd.h></code>, this value is the set of libraries to be given to the <code>cc</code> and <code>c89</code> utilities to build an application using the Large File Summit transitional compilation environment (see <code>lfcompile64(5)</code>).<code>_CS_LFS64_LINTFLAGS</code> If <code>_LFS64_LARGEFILE</code> is defined in <code><unistd.h></code>, this value is the set of options to be given to the <code>lint</code> utility to check application source using the Large File Summit transitional compilation environment (see <code>lfcompile64(5)</code>).<code>_CS_LFS_CFLAGS</code> If <code>_LFS_LARGEFILE</code> is defined in <code><unistd.h></code>, this value is the set of initial options to be given to the <code>cc</code> and <code>c89</code> utilities to build an application using the Large File Summit large file compilation environment for 32-bit applications (see <code>lfcompile(5)</code>).

_CS_LFS_LDFLAGS

If `_LFS_LARGEFILE` is defined in `<unistd.h>`, this value is the set of final options to be given to the `cc` and `c89` utilities to build an application using the Large File Summit large file compilation environment for 32-bit applications (see `lfcompile(5)`).

_CS_LFS_LIBS

If `_LFS_LARGEFILE` is defined in `<unistd.h>`, this value is the set of libraries to be given to the `cc` and `c89` utilities to build an application using the Large File Summit large file compilation environment for 32-bit applications (see `lfcompile(5)`).

_CS_LFS_LINTFLAGS

If `_LFS_LARGEFILE` is defined in `<unistd.h>`, this value is the set of options to be given to the `lint` utility to check application source using the Large File Summit large file compilation environment for 32-bit applications (see `lfcompile(5)`).

_CS_PATH

If the ISO POSIX.2 standard is supported, this is the value for the `PATH` environment variable that finds all standard utilities. Otherwise the meaning of this value is unspecified.

_CS_POSIX_V6_ILP32_OFF32_CFLAGS

If `sysconf(_SC_V6_ILP32_OFF32)` returns -1, the meaning of this value is unspecified. Otherwise, this value is the set of initial options to be given to the `c99` utility to build an application using a programming model with 32-bit `int`, `long`, `pointer`, and `off_t` types.

_CS_POSIX_V6_ILP32_OFF32_LDFLAGS

If `sysconf(_SC_V6_ILP32_OFF32)` returns -1, the meaning of this value is unspecified. Otherwise, this value is the set of final options to be given to the `c99` utility to build an application using a programming model with 32-bit `int`, `long`, `pointer`, and `off_t` types.

_CS_POSIX_V6_ILP32_OFF32_LIBS

If `sysconf(_SC_V6_ILP32_OFF32)` returns -1, the meaning of this value is unspecified. Otherwise, this value is the set of libraries to be given to the `c99` utility to build an application using a programming model with 32-bit `int`, `long`, `pointer`, and `off_t` types.

_CS_POSIX_V6_ILP32_OFFBIG_CFLAGS

If `sysconf(_SC_V6_ILP32_OFFBIG)` returns -1, the meaning of this value is unspecified. Otherwise, this value is the set of initial options to be given to the `c99` utility to build an application using a programming model with 32-bit `int`, `long`, and `pointer` types, and an `off_t` type using at least 64 bits.

_CS_POSIX_V6_ILP32_OFFBIG_LDFLAGS

If `sysconf(_SC_V6_ILP32_OFFBIG)` returns -1, the meaning of this value is unspecified. Otherwise, this value is the set of final options to be given to the `c99` utility to build an application using a programming model with 32-bit `int`, `long`, and `pointer` types, and an `off_t` type using at least 64 bits.

confstr(3C)

`_CS_POSIX_V6_ILP32_OFFBIG_LIBS`

If `sysconf(_SC_V6_ILP32_OFFBIG)` returns -1, the meaning of this value is unspecified. Otherwise, this value is the set of libraries to be given to the `c99` utility to build an application using a programming model with 32-bit `int`, `long`, and `pointer` types, and an `off_t` type using at least 64 bits.

`_CS_POSIX_V6_LP64_OFF64_CFLAGS`

If `sysconf(_SC_V6_LP64_OFF64)` returns -1, the meaning of this value is unspecified. Otherwise, this value is the set of initial options to be given to the `c99` utility to build an application using a programming model with 64-bit `int`, `long`, `pointer`, and `off_t` types.

`_CS_POSIX_V6_LP64_OFF64_LDFLAGS`

If `sysconf(_SC_V6_LP64_OFF64)` returns -1, the meaning of this value is unspecified. Otherwise, this value is the set of final options to be given to the `c99` utility to build an application using a programming model with 64-bit `int`, `long`, `pointer`, and `off_t` types.

`_CS_POSIX_V6_LP64_OFF64_LIBS`

If `sysconf(_SC_V6_LP64_OFF64)` returns -1, the meaning of this value is unspecified. Otherwise, this value is the set of libraries to be given to the `c99` utility to build an application using a programming model with 64-bit `int`, `long`, `pointer`, and `off_t` types.

`_CS_POSIX_V6_LPBIG_OFFBIG_CFLAGS`

If `sysconf(_SC_V6_LPBIG_OFFBIG)` returns -1, the meaning of this value is unspecified. Otherwise, this value is the set of initial options to be given to the `c99` utility to build an application using a programming model with an `int` type using at least 32 bits and `long`, `pointer`, and `off_t` types using at least 64 bits.

`_CS_POSIX_V6_LPBIG_OFFBIG_LDFLAGS`

If `sysconf(_SC_V6_LPBIG_OFFBIG)` returns -1, the meaning of this value is unspecified. Otherwise, this value is the set of final options to be given to the `c99` utility to build an application using a programming model with an `int` type using at least 32 bits and `long`, `pointer`, and `off_t` types using at least 64 bits.

`_CS_POSIX_V6_LPBIG_OFFBIG_LIBS`

If `sysconf(_SC_V6_LPBIG_OFFBIG)` returns -1, the meaning of this value is unspecified. Otherwise, this value is the set of libraries to be given to the `c99` utility to build an application using a programming model with an `int` type using at least 32 bits and `long`, `pointer`, and `off_t` types using at least 64 bits.

`_CS_POSIX_V6_WIDTH_RESTRICTED_ENVS`

This value is a <newline>-separated list of names of programming environments supported by the implementation in which the widths of the `blksize_t`, `cc_t`, `mode_t`, `nfds_t`, `pid_t`, `ptrdiff_t`, `size_t`, `speed_t`, `ssize_t`, `suseconds_t`, `tcflag_t`, `useconds_t`, `wchar_t`, and `wint_t` types are no greater than the width of type `long`.

`_CS_XBS5_ILP32_OFF32_CFLAGS`

If `sysconf(_SC_XBS5_ILP32_OFF32)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of initial options to be given to the `cc` and `c89` utilities to build an application using a programming model with 32-bit `int`, `long`, `pointer`, and `off_t` types.

`_CS_XBS5_ILP32_OFF32_LDFLAGS`

If `sysconf(_SC_XBS5_ILP32_OFF32)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of final options to be given to the `cc` and `c89` utilities to build an application using a programming model with 32-bit `int`, `long`, `pointer`, and `off_t` types.

`_CS_XBS5_ILP32_OFF32_LIBS`

If `sysconf(_SC_XBS5_ILP32_OFF32)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of libraries to be given to the `cc` and `c89` utilities to build an application using a programming model with 32-bit `int`, `long`, `pointer`, and `off_t` types.

`_CS_XBS5_ILP32_OFF32_LINTFLAGS`

If `sysconf(_SC_XBS5_ILP32_OFF32)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of options to be given to the `lint` utility to check application source using a programming model with 32-bit `int`, `long`, `pointer`, and `off_t` types.

`_CS_XBS5_ILP32_OFFBIG_CFLAGS`

If `sysconf(_SC_XBS5_ILP32_OFFBIG)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of initial options to be given to the `cc` and `c89` utilities to build an application using a programming model with 32-bit `int`, `long`, and `pointer` types, and an `off_t` type using at least 64 bits.

`_CS_XBS5_ILP32_OFFBIG_LDFLAGS`

If `sysconf(_SC_XBS5_ILP32_OFFBIG)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of final options to be given to the `cc` and `c89` utilities to build an application using a programming model with 32-bit `int`, `long`, and `pointer` types, and an `off_t` type using at least 64 bits.

`_CS_XBS5_ILP32_OFFBIG_LIBS`

If `sysconf(_SC_XBS5_ILP32_OFFBIG)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of libraries to be given to the `cc` and `c89` utilities to build an application using a programming model with 32-bit `int`, `long`, and `pointer` types, and an `off_t` type using at least 64 bits.

`_CS_XBS5_ILP32_OFFBIG_LINTFLAGS`

If `sysconf(_SC_XBS5_ILP32_OFFBIG)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of options to be given to the `lint` utility to check an application using a programming model with 32-bit `int`, `long`, and `pointer` types, and an `off_t` type using at least 64 bits.

The `confstr()` function supports the following values for *name*, defined in `<unistd.h>`, for SPARC only:

confstr(3C)

`_CS_XBS5_LP64_OFF64_CFLAGS`

If `sysconf(_SC_XBS5_LP64_OFF64)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of initial options to be given to the `cc` and `c89` utilities to build an application using a programming model with 64-bit `int`, `long`, `pointer`, and `off_t` types.

`_CS_XBS5_LP64_OFF64_LDFLAGS`

If `sysconf(_SC_XBS5_LP64_OFF64)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of final options to be given to the `cc` and `c89` utilities to build an application using a programming model with 64-bit `int`, `long`, `pointer`, and `off_t` types.

`_CS_XBS5_LP64_OFF64_LIBS`

If `sysconf(_SC_XBS5_LP64_OFF64)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of libraries to be given to the `cc` and `c89` utilities to build an application using a programming model with 64-bit `int`, `long`, `pointer`, and `off_t` types.

`_CS_XBS5_LP64_OFF64_LINTFLAGS`

If `sysconf(_SC_XBS5_LP64_OFF64)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of options to be given to the `lint` utility to check application source using a programming model with 64-bit `int`, `long`, `pointer`, and `off_t` types.

`_CS_XBS5_LPBIG_OFFBIG_CFLAGS`

If `sysconf(_SC_XBS5_LPBIG_OFFBIG)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of initial options to be given to the `cc` and `c89` utilities to build an application using a programming model with an `int` type using at least 32 bits and `long`, `pointer`, and `off_t` types using at least 64 bits.

`_CS_XBS5_LPBIG_OFFBIG_LDFLAGS`

If `sysconf(_SC_XBS5_LPBIG_OFFBIG)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of final options to be given to the `cc` and `c89` utilities to build an application using a programming model with an `int` type using at least 32 bits and `long`, `pointer`, and `off_t` types using at least 64 bits.

`_CS_XBS5_LPBIG_OFFBIG_LIBS`

If `sysconf(_SC_XBS5_LPBIG_OFFBIG)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of libraries to be given to the `cc` and `c89` utilities to build an application using a programming model with an `int` type using at least 32 bits and `long`, `pointer`, and `off_t` types using at least 64 bits.

`_CS_XBS5_LPBIG_OFFBIG_LINTFLAGS`

If `sysconf(_SC_XBS5_LPBIG_OFFBIG)` returns `-1` the meaning of this value is unspecified. Otherwise, this value is the set of options to be given to the `lint` utility to check application source using a programming model with an `int` type using at least 32 bits and `long`, `pointer`, and `off_t` types using at least 64 bits.

RETURN VALUES If *name* has a configuration-defined value, the `confstr()` function returns the size of buffer that would be needed to hold the entire configuration-defined value. If this return value is greater than *len*, the string returned in *buf* is truncated.

If *name* is invalid, `confstr()` returns 0 and sets `errno` to indicate the error.

If *name* does not have a configuration-defined value, `confstr()` returns 0 and leaves `errno` unchanged.

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

`EINVAL` The value of the *name* argument is invalid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Mt-Safe

SEE ALSO `pathconf(2)`, `sysconf(3C)`, `attributes(5)`, `lfcompile(5)`, `lfcompile64(5)`, `standards(5)`

crypt(3C)

NAME	crypt – string encoding function								
SYNOPSIS	<pre>#include <crypt.h> char *crypt(const char *key, const char *salt);</pre>								
Standard conforming	<pre>#include <unistd.h> char *crypt(const char *key, const char *salt);</pre>								
DESCRIPTION	<p>The <code>crypt()</code> function encodes strings suitable for secure storage as passwords. It generates the password hash given the <i>key</i> and <i>salt</i>.</p> <p>The <i>key</i> argument is the plain text password to be encrypted.</p> <p>If the first character of <i>salt</i> is "\$", <code>crypt()</code> uses <code>crypt.conf(4)</code> to determine which shared module to load for the encryption algorithm. The algorithm name <code>crypt()</code> uses to search in <code>crypt.conf</code> is the string between the first and second "\$", or between the first "\$" and first "," if a "," comes before the second "\$".</p> <p>If the first character of <i>salt</i> is not "\$", the algorithm described on <code>crypt_unix(5)</code> is used.</p>								
RETURN VALUES	<p>Upon successful completion, <code>crypt()</code> returns a pointer to the encoded string. Otherwise it returns a null pointer and sets <code>errno</code> to indicate the error.</p> <p>The return value points to static data that is overwritten by each call.</p>								
ERRORS	<p>The <code>crypt()</code> function will fail if:</p> <table><tr><td>EINVAL</td><td>An entry in <code>crypt.conf</code> is invalid.</td></tr><tr><td>ELIBACC</td><td>The required shared library was not found.</td></tr><tr><td>ENOMEM</td><td>There is insufficient memory to generate the hash.</td></tr><tr><td>ENOSYS</td><td>The functionality is not supported on this system.</td></tr></table>	EINVAL	An entry in <code>crypt.conf</code> is invalid.	ELIBACC	The required shared library was not found.	ENOMEM	There is insufficient memory to generate the hash.	ENOSYS	The functionality is not supported on this system.
EINVAL	An entry in <code>crypt.conf</code> is invalid.								
ELIBACC	The required shared library was not found.								
ENOMEM	There is insufficient memory to generate the hash.								
ENOSYS	The functionality is not supported on this system.								
USAGE	<p>The values returned by this function might not be portable among standard-conforming systems. See <code>standards(5)</code>.</p> <p>Applications should not use <code>crypt()</code> to store or verify user passwords but should use the functions described on <code>pam(3PAM)</code> instead.</p>								
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:								

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

crypt(3C)

SEE ALSO | passwd(1), crypt_genhash_impl(3C), crypt_gensalt(3C),
crypt_gensalt_impl(3C), getpassphrase(3C), pam(3PAM), passwd(4),
policy.conf(4), attributes(5), crypt_unix(5), standards(5)

crypt_genhash_impl(3C)

NAME	crypt_genhash_impl – generate encrypted password						
SYNOPSIS	<pre>#include <crypt.h> char *crypt_genhash_impl(char *ctbuffer, size_t ctbufferlen, const char *plaintext, const char *salt, const char **params);</pre>						
DESCRIPTION	<p>The <code>crypt_genhash_impl()</code> function is called by crypt(3C) to generate the encrypted password <i>plaintext</i>.</p> <p>The <i>ctbuffer</i> argument is a pointer to an MT-safe buffer of <i>ctbufferlen</i> size that is used to return the result.</p> <p>The <i>salt</i> argument is the salt used in encoding.</p> <p>The <i>params</i> argument is an <i>argv</i>-like null-terminated vector of type <code>char *</code>. The first element of <i>params</i> represents the mechanism token name from <code>crypt.conf(4)</code>. The remaining elements of <i>params</i> represent strings of the form <code><parameter>[=<value>]</code> to allow passing in additional information from the <code>crypt.conf</code> entry, such as specifying rounds information "rounds=4096".</p> <p>The <code>crypt_genhash_impl()</code> function must not free(3C) <i>ctbufferlen</i> on error.</p>						
RETURN VALUES	Upon successful completion, <code>crypt_genhash_impl()</code> returns a pointer to the encoded version of <i>plaintext</i> . Otherwise a null pointer is returned and <code>errno</code> is set to indicate the error.						
ERRORS	The <code>crypt_genhash_impl()</code> function will fail if: <table><tr><td>EINVAL</td><td>The configuration file <code>crypt.conf</code> contains an invalid entry.</td></tr><tr><td>ELIBACC</td><td>The required shared library was not found.</td></tr><tr><td>ENOMEM</td><td>There is insufficient memory to perform hashing.</td></tr></table>	EINVAL	The configuration file <code>crypt.conf</code> contains an invalid entry.	ELIBACC	The required shared library was not found.	ENOMEM	There is insufficient memory to perform hashing.
EINVAL	The configuration file <code>crypt.conf</code> contains an invalid entry.						
ELIBACC	The required shared library was not found.						
ENOMEM	There is insufficient memory to perform hashing.						
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>Interface Stability</td><td>Evolving</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
MT-Level	MT-Safe						
SEE ALSO	passwd(1) , crypt(3C) , crypt_gensalt_impl(3C) , free(3C) , getpassphrase(3C) , crypt.conf(4) , passwd(4) , attributes(5)						

NAME	crypt_gensalt – generate salt string for string encoding						
SYNOPSIS	<pre>#include <crypt.h> char *crypt_gensalt(const char *oldsalt, const struct passwd *userinfo);</pre>						
DESCRIPTION	<p>The <code>crypt_gensalt()</code> function generates the salt string required by <code>crypt(3C)</code>.</p> <p>If <code>oldsalt</code> is <code>NULL</code>, <code>crypt_gensalt()</code> uses the algorithm defined by <code>CRYPT_DEFAULT</code> in <code>/etc/security/policy.conf</code>. See <code>policy.conf(4)</code>.</p> <p>If <code>oldsalt</code> is non-null, <code>crypt_gensalt()</code> determines if the algorithm specified by <code>oldsalt</code> is allowable by checking the <code>CRYPT_ALGORITHMS_ALLOW</code> and <code>CRYPT_ALGORITHMS_DEPRECATED</code> variables in <code>/etc/security/policy.conf</code>. If the algorithm is allowed, <code>crypt_gensalt()</code> loads the appropriate shared library and calls <code>crypt_gensalt_impl(3C)</code>. If the algorithm is not allowed or there is no entry for it in <code>crypt.conf</code>, <code>crypt_gensalt()</code> uses the default algorithm.</p> <p>The mechanism just described provides a means to migrate users to new password hashing algorithms when the password is changed.</p>						
RETURN VALUES	Upon successful completion, <code>crypt_gensalt()</code> returns a pointer to the new salt. Otherwise a null pointer is returned and <code>errno</code> is set to indicate the error.						
ERRORS	<p>The <code>crypt_gensalt()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>EINVAL</code></td> <td>The configuration file <code>crypt.conf</code> contains an invalid entry.</td> </tr> <tr> <td><code>ELIBACC</code></td> <td>The required shared library was not found.</td> </tr> <tr> <td><code>ENOMEM</code></td> <td>There is insufficient memory to perform hashing.</td> </tr> </table>	<code>EINVAL</code>	The configuration file <code>crypt.conf</code> contains an invalid entry.	<code>ELIBACC</code>	The required shared library was not found.	<code>ENOMEM</code>	There is insufficient memory to perform hashing.
<code>EINVAL</code>	The configuration file <code>crypt.conf</code> contains an invalid entry.						
<code>ELIBACC</code>	The required shared library was not found.						
<code>ENOMEM</code>	There is insufficient memory to perform hashing.						
USAGE	<p>The value returned by <code>crypt_gensalt()</code> points to a null-terminated string. The caller of <code>crypt_gensalt()</code> is responsible for calling <code>free(3C)</code>.</p> <p>Applications dealing with user authentication and password changing should not call <code>crypt_gensalt()</code> directly but should instead call the appropriate <code>pam(3PAM)</code> functions.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<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>Interface Stability</td> <td>Evolving</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
MT-Level	MT-Safe						
SEE ALSO	<code>passwd(1)</code> , <code>crypt(3C)</code> , <code>crypt_genhash_impl(3C)</code> , <code>crypt_gensalt_impl(3C)</code> , <code>getpassphrase(3C)</code> , <code>malloc(3C)</code> , <code>pam(3PAM)</code> , <code>crypt.conf(4)</code> , <code>passwd(4)</code> , <code>policy.conf(4)</code> , <code>attributes(5)</code>						

crypt_gensalt_impl(3C)

NAME	crypt_gensalt_impl – generate salt for password encryption						
SYNOPSIS	<pre>#include <crypt.h> char *crypt_gensalt_impl(char *gsbuffer, size_t gsbuflen, const char *oldsalt, const struct passwd *userinfo, const char **params);</pre>						
DESCRIPTION	<p>The <code>crypt_gensalt_impl()</code> function is called by <code>crypt_gensalt(3C)</code> to generate the salt for password encryption.</p> <p>The <code>gsbuffer</code> argument is a pointer to an MT-safe buffer of size <code>gsbuflen</code>.</p> <p>The <code>oldsalt</code> and <code>userinfo</code> arguments are passed unchanged from <code>crypt_gensalt(3C)</code>.</p> <p>The <code>params</code> argument is an <i>argv</i>-like null terminated vector of type <code>char *</code>. The first element of <code>params</code> represents the mechanism token name from <code>crypt.conf(4)</code>. The remaining elements of <code>params</code> represent strings of the form <code><parameter>[=<value>]</code> to allow passing in additional information from the <code>crypt.conf</code> entry, such as specifying rounds information "rounds=4096".</p> <p>The value returned by <code>crypt_gensalt_impl()</code> points to a thread-specific buffer to be freed by the caller of <code>crypt_gensalt(3C)</code> after calling <code>crypt(3C)</code>.</p>						
RETURN VALUES	Upon successful completion, <code>crypt_gensalt_impl()</code> returns a pointer to the new salt. Otherwise a null pointer is returned and <code>errno</code> is set to indicate the error.						
ERRORS	The <code>crypt_gensalt_impl()</code> function will fail if: <table><tr><td>EINVAL</td><td>The configuration file <code>crypt.conf</code> contains an invalid entry.</td></tr><tr><td>ELIBACC</td><td>The required crypt shared library was not found.</td></tr><tr><td>ENOMEM</td><td>There is insufficient memory to perform hashing.</td></tr></table>	EINVAL	The configuration file <code>crypt.conf</code> contains an invalid entry.	ELIBACC	The required crypt shared library was not found.	ENOMEM	There is insufficient memory to perform hashing.
EINVAL	The configuration file <code>crypt.conf</code> contains an invalid entry.						
ELIBACC	The required crypt shared library was not found.						
ENOMEM	There is insufficient memory to perform hashing.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes: <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Evolving</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
MT-Level	MT-Safe						
SEE ALSO	<code>passwd(1)</code> , <code>crypt(3C)</code> , <code>crypt_genhash_impl(3C)</code> , <code>crypt_gensalt(3C)</code> , <code>getpassphrase(3C)</code> , <code>crypt.conf(4)</code> , <code>passwd(4)</code> , <code>attributes(5)</code>						

NAME	cset, csetlen, csetcol, csetno, wcsetno – get information on EUC codesets				
SYNOPSIS	<pre>#include <euc.h> int csetlen(int codeset); int csetcol(int codeset); int csetno(unsigned char c); #include <widec.h> int wcsetno(wchar_t pc);</pre>				
DESCRIPTION	<p>Both <code>csetlen()</code> and <code>csetcol()</code> take a code set number <i>codeset</i>, which must be 0, 1, 2, or 3. The <code>csetlen()</code> function returns the number of bytes needed to represent a character of the given Extended Unix Code (EUC) code set, excluding the single-shift characters SS2 and SS3 for codesets 2 and 3. The <code>csetcol()</code> function returns the number of columns a character in the given EUC code set would take on the display.</p> <p>The <code>csetno()</code> function is implemented as a macro that returns a codeset number (0, 1, 2, or 3) for the EUC character whose first byte is <i>c</i>. For example,</p> <pre>#include<euc.h> . . . x+=csetcol(csetno(c));</pre> <p>increments a counter “<i>x</i>” (such as the cursor position) by the width of the character whose first byte is <i>c</i>.</p> <p>The <code>wcsetno()</code> function is implemented as a macro that returns a codeset number (0, 1, 2, or 3) for the given process code character <i>pc</i>. For example,</p> <pre>#include<euc.h> #include<widec.h> . . . x+=csetcol(wcsetno(pc));</pre> <p>increments a counter “<i>x</i>” (such as the cursor position) by the width of the Process Code character <i>pc</i>.</p>				
USAGE	<p>These functions work only for the EUC locales.</p> <p>The <code>cset()</code>, <code>csetlen()</code>, <code>csetcol()</code>, <code>csetno()</code>, and <code>wcsetno()</code> functions can be used safely in multithreaded applications, as long as <code>setlocale(3C)</code> is not being called to change the locale.</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<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>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe with exceptions				
SEE ALSO	<code>setlocale(3C)</code> , <code>euclen(3C)</code> , <code>attributes(5)</code>				

ctermid(3C)

NAME	ctermid, ctermid_r – generate path name for controlling terminal						
SYNOPSIS	<pre>#include <stdio.h> char *ctermid(char *s); char *ctermid_r(char *s);</pre>						
DESCRIPTION							
ctermid()	<p>The <code>ctermid()</code> function generates the path name of the controlling terminal for the current process and stores it in a string.</p> <p>If <code>s</code> is a null pointer, the string is stored in an internal static area whose address is returned and whose contents are overwritten at the next call to <code>ctermid()</code>. Otherwise, <code>s</code> is assumed to point to a character array of at least <code>L_ctermid</code> elements. The path name is placed in this array and the value of <code>s</code> is returned. The constant <code>L_ctermid</code> is defined in the header <code><stdio.h></code>.</p>						
ctermid_r()	<p>The <code>ctermid_r()</code> function behaves as <code>ctermid()</code> except that if <code>s</code> is a null pointer, the function returns <code>NULL</code>.</p>						
USAGE	<p>The difference between <code>ctermid()</code> and ttyname(3C) is that <code>ttyname()</code> must be passed a file descriptor and returns the actual name of the terminal associated with that file descriptor, while <code>ctermid()</code> returns a string (<code>/dev/tty</code>) that will refer to the terminal if used as a file name. The <code>ttyname()</code> function is useful only if the process already has at least one file open to a terminal.</p> <p>The <code>ctermid()</code> function is unsafe in multithreaded applications. The <code>ctermid_r()</code> function is MT-Safe and should be used instead.</p> <p>When compiling multithreaded applications, the <code>_REENTRANT</code> flag must be defined on the compile line. This flag should be used only with multithreaded applications.</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>Interface Stability</td><td><code>ctermid()</code> is Standard</td></tr><tr><td>MT-Level</td><td><code>ctermid()</code> is Unsafe; <code>ctermid_r()</code> is MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	<code>ctermid()</code> is Standard	MT-Level	<code>ctermid()</code> is Unsafe; <code>ctermid_r()</code> is MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	<code>ctermid()</code> is Standard						
MT-Level	<code>ctermid()</code> is Unsafe; <code>ctermid_r()</code> is MT-Safe						
SEE ALSO	ttyname(3C) , attributes(5)						

NAME	ctime, ctime_r, localtime, localtime_r, gmtime, gmtime_r, asctime, asctime_r, tzset – convert date and time to string
SYNOPSIS	<pre>#include <time.h> char *ctime(const time_t *clock); struct tm *localtime(const time_t *clock); struct tm *gmtime(const time_t *clock); char *asctime(const struct tm *tm); extern time_t timezone, altzone; extern int daylight; extern char *tzname[2]; void tzset(void); char *ctime_r(const time_t *clock, char *buf, int buflen); struct tm *localtime_r(const time_t *restrict clock, struct tm *restrict res); struct tm *gmtime_r(const time_t *restrict clock, struct tm *restrict res); char *asctime_r(const struct tm *restrict tm, char *restrict buf, int buflen);</pre>
Standard conforming	<pre>cc [flag...] file... -D_POSIX_PTHREAD_SEMANTICS [library...] char *ctime_r(const time_t *clock, char *buf); char *asctime_r(const struct tm *tm, char *buf);</pre>
DESCRIPTION	<p>The <code>ctime()</code> function converts the time pointed to by <code>clock</code>, representing the time in seconds since the Epoch (00:00:00 UTC, January 1, 1970), to local time in the form of a 26-character string, as shown below. Time zone and daylight savings corrections are made before string generation. The fields are in constant width:</p> <pre>Fri Sep 13 00:00:00 1986\n\0</pre> <p>The <code>ctime()</code> function is equivalent to:</p> <pre>asctime(localtime(clock))</pre> <p>The <code>ctime()</code>, <code>asctime()</code>, <code>gmtime()</code>, and <code>localtime()</code> functions return values in one of two thread-specific data objects: a broken-down time structure and an array of <code>char</code>. Execution of any of the functions can overwrite the information returned in either of these objects by any of the other functions executed by the same thread.</p>

ctime(3C)

The `ctime_r()` function has the same functionality as `ctime()` except that the caller must supply a buffer *buf* with length *buflen* to store the result; *buf* must be at least 26 bytes. The standard-conforming `ctime_r()` function does not take a *buflen* parameter.

The `localtime()` and `gmtime()` functions return pointers to `tm` structures (see below). The `localtime()` function corrects for the main time zone and possible alternate ("daylight savings") time zone; the `gmtime()` function converts directly to Coordinated Universal Time (UTC), which is what the UNIX system uses internally.

The `localtime_r()` and `gmtime_r()` functions have the same functionality as `localtime()` and `gmtime()` respectively, except that the caller must supply a buffer *res* to store the result.

The `asctime()` function converts a `tm` structure to a 26-character string, as shown in the previous example, and returns a pointer to the string.

The `asctime_r()` function has the same functionality as `asctime()` except that the caller must supply a buffer *buf* with length *buflen* for the result to be stored. The *buf* argument must be at least 26 bytes. The standard-conforming `asctime_r()` function does not take a *buflen* parameter. The `asctime_r()` function returns a pointer to *buf* upon success. In case of failure, `NULL` is returned and `errno` is set.

Declarations of all the functions and externals, and the `tm` structure, are in the `<time.h>` header. The members of the `tm` structure are:

```
int    tm_sec;    /* seconds after the minute - [0, 60] */
                /* for leap seconds */
int    tm_min;    /* minutes after the hour - [0, 59] */
int    tm_hour;   /* hour since midnight - [0, 23] */
int    tm_mday;   /* day of the month - [1, 31] */
int    tm_mon;    /* months since January - [0, 11] */
int    tm_year;   /* years since 1900 */
int    tm_wday;   /* days since Sunday - [0, 6] */
int    tm_yday;   /* days since January 1 - [0, 365] */
int    tm_isdst;  /* flag for alternate daylight savings time */
```

The value of `tm_isdst` is positive if daylight savings time is in effect, zero if daylight savings time is not in effect, and negative if the information is not available. Previously, the value of `tm_isdst` was defined as non-zero if daylight savings was in effect.

The external `time_t` variable `altzone` contains the difference, in seconds, between Coordinated Universal Time and the alternate time zone. The external variable `timezone` contains the difference, in seconds, between UTC and local standard time. The external variable `daylight` indicates whether time should reflect daylight savings time. Both `timezone` and `altzone` default to 0 (UTC). The external variable `daylight` is non-zero if an alternate time zone exists. The time zone names are contained in the external variable `tzname`, which by default is set to:

```
char *tzname[2] = { "GMT", "" };
```


These functions know about the peculiarities of this conversion for various time periods for the U.S. (specifically, the years 1974, 1975, and 1987). They start handling the new daylight savings time starting with the first Sunday in April, 1987.

The `tzset()` function uses the contents of the environment variable `TZ` to override the value of the different external variables. It is called by `asctime()` and can also be called by the user. See `environ(5)` for a description of the `TZ` environment variable.

Starting and ending times are relative to the current local time zone. If the alternate time zone start and end dates and the time are not provided, the days for the United States that year will be used and the time will be 2 AM. If the start and end dates are provided but the time is not provided, the time will be 2 AM. The effects of `tzset()` change the values of the external variables `timezone`, `altzone`, `daylight`, and `tzname`.

Note that in most installations, `TZ` is set to the correct value by default when the user logs on, using the local `/etc/default/init` file (see `TIMEZONE(4)`).

RETURN VALUES

Upon successful completion, the `gmtime()` and `localtime()` functions return a pointer to a `struct tm`. If an error is detected, `gmtime()` and `localtime()` return a null pointer.

Upon successful completion, the `gmtime_r()` and `localtime_r()` functions return the address of the structure pointed to by the `res` argument. If an error is detected, `gmtime_r()` and `localtime_r()` return a null pointer and set `errno` to indicate the error.

ERRORS

The `ctime_r()` and `asctime_r()` functions will fail if:

ERANGE The length of the buffer supplied by the caller is not large enough to store the result.

The `gmtime()`, `gmtime_r()`, `localtime()`, and `localtime_r()` functions will fail if:

EOVERFLOW The result cannot be represented.

USAGE

These functions do not support localized date and time formats. The `strftime(3C)` function can be used when localization is required.

The `localtime()`, `localtime_r()`, `gmtime()`, `gmtime_r()`, `ctime()`, and `ctime_r()` functions assume Gregorian dates. Times before the adoption of the Gregorian calendar will not match historical records.

EXAMPLES

EXAMPLE 1 Examples of the `tzset()` function.

The `tzset()` function scans the contents of the environment variable and assigns the different fields to the respective variable. For example, the most complete setting for New Jersey in 1986 could be:

```
EST5EDT4,116/2:00:00,298/2:00:00
or simply
```

ctime(3C)

EXAMPLE 1 Examples of the `tzset()` function. (Continued)

```
EST5EDT
```

An example of a southern hemisphere setting such as the Cook Islands could be

```
KDT9:30KST10:00,63/5:00,302/20:00
```

In the longer version of the New Jersey example of TZ, `tzname[0]` is EST, `timezone` is set to $5*60*60$, `tzname[1]` is EDT, `altzone` is set to $4*60*60$, the starting date of the alternate time zone is the 117th day at 2 AM, the ending date of the alternate time zone is the 299th day at 2 AM (using zero-based Julian days), and `daylight` is set positive. Starting and ending times are relative to the current local time zone. If the alternate time zone start and end dates and the time are not provided, the days for the United States that year will be used and the time will be 2 AM. If the start and end dates are provided but the time is not provided, the time will be 2 AM. The effects of `tzset()` are thus to change the values of the external variables `timezone`, `altzone`, `daylight`, and `tzname`. The `ctime()`, `localtime()`, `mkttime()`, and `strftime()` functions also update these external variables as if they had called `tzset()` at the time specified by the `time_t` or `struct tm` value that they are converting.

BUGS The `zoneinfo` `timezone` data files do not transition past Tue Jan 19 03:14:07 2038 UTC. Therefore for 64-bit applications using `zoneinfo` timezones, calculations beyond this date might not use the correct offset from standard time, and could return incorrect values. This affects the 64-bit version of `localtime()`, `localtime_r()`, `ctime()`, and `ctime_r()`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

The `asctime()`, `ctime()`, `gmtime()`, and `localtime()` functions are safe to use in multithread applications because they employ thread-specific data. However, their use is discouraged because standards do not require them to be thread-safe. The `asctime_r()` and `gmtime_r()` functions are MT-Safe. The `ctime_r()`, `localtime_r()`, and `tzset()` functions are MT-Safe in multithread applications, as long as no user-defined function directly modifies one of the following variables: `timezone`, `altzone`, `daylight`, and `tzname`. These four variables are not MT-Safe to access. They are modified by the `tzset()` function in an MT-Safe manner. The `mktime()`, `localtime_r()`, and `ctime_r()` functions call `tzset()`.

SEE ALSO time(2), Intro(3), [getenv\(3C\)](#), [mktime\(3C\)](#), [printf\(3C\)](#), [putenv\(3C\)](#), [setlocale\(3C\)](#), [strftime\(3C\)](#), [TIMEZONE\(4\)](#), [attributes\(5\)](#), [environ\(5\)](#), [standards\(5\)](#)

NOTES When compiling multithreaded programs, see [Intro\(3\)](#).

The return values for `asctime()`, `ctime()`, `gmtime()`, and `localtime()` point to thread-specific data whose content is overwritten by each call by the same thread.

Setting the time during the interval of change from `timezone` to `altzone` or vice versa can produce unpredictable results. The system administrator must change the Julian start and end days annually.

Solaris 2.4 and earlier releases provided definitions of the `ctime_r()`, `localtime_r()`, `gmtime_r()`, and `asctime_r()` functions as specified in POSIX.1c Draft 6. The final POSIX.1c standard changed the interface for `ctime_r()` and `asctime_r()`. Support for the Draft 6 interface is provided for compatibility only and might not be supported in future releases. New applications and libraries should use the standard-conforming interface.

For POSIX.1c-conforming applications, the `_POSIX_PTHREAD_SEMANTICS` and `_REENTRANT` flags are automatically turned on by defining the `_POSIX_C_SOURCE` flag with a value $\geq 199506L$.

In Solaris 10, `gmtime()`, `gmtime_r()`, `localtime()`, and `localtime_r()` were updated to return a null pointer if an error is detected. This change was based on the SUSv3 specification. See [standards\(5\)](#).

ctype(3C)

NAME	ctype, isalpha, isalnum, isascii, isblank, iscntrl, isdigit, islower, isprint, isspace, isupper, ispunct, isgraph, isxdigit – character handling		
SYNOPSIS	<pre>#include <ctype.h> int isalpha(int c); int isalnum(int c); int isascii(int c); int isblank(int c); int iscntrl(int c); int isdigit(int c); int isgraph(int c); int islower(int c); int isprint(int c); int ispunct(int c); int isspace(int c); int isupper(int c); int isxdigit(int c);</pre>		
DESCRIPTION	<p>These macros classify character-coded integer values. Each is a predicate returning non-zero for true, 0 for false. The behavior of these macros, except <code>isascii()</code>, is affected by the current locale (see setlocale(3C)). To modify the behavior, change the <code>LC_TYPE</code> category in <code>setlocale()</code>, that is, <code>setlocale(LC_TYPE, newlocale)</code>. In the "C" locale, or in a locale where character type information is not defined, characters are classified according to the rules of the US-ASCII 7-bit coded character set.</p> <p>The <code>isascii()</code> macro is defined on all integer values. The rest are defined only where the argument is an <code>int</code>, the value of which is representable as an unsigned <code>char</code>, or <code>EOF</code>, which is defined by the <code><stdio.h></code> header and represents end-of-file.</p> <p>Functions exist for all the macros defined below. To get the function form, the macro name must be undefined (for example, <code>#undef isdigit</code>).</p> <p>For macros described with <code>Default</code> and <code>Standard conforming</code> versions, standard-conforming behavior is provided for standard-conforming applications (see standards(5)) and for applications that define <code>__XPG4_CHAR_CLASS__</code> before including <code><ctype.h></code>.</p>		
Default	<table><tr><td><code>isalpha()</code></td><td>Tests for any character for which <code>isupper()</code> or <code>islower()</code> is true.</td></tr></table>	<code>isalpha()</code>	Tests for any character for which <code>isupper()</code> or <code>islower()</code> is true.
<code>isalpha()</code>	Tests for any character for which <code>isupper()</code> or <code>islower()</code> is true.		
Standard conforming	<table><tr><td><code>isalpha()</code></td><td>Tests for any character for which <code>isupper()</code> or <code>islower()</code> is true, or any character that is one of the current locale-defined set of</td></tr></table>	<code>isalpha()</code>	Tests for any character for which <code>isupper()</code> or <code>islower()</code> is true, or any character that is one of the current locale-defined set of
<code>isalpha()</code>	Tests for any character for which <code>isupper()</code> or <code>islower()</code> is true, or any character that is one of the current locale-defined set of		

		characters for which none of <code>iscntrl()</code> , <code>isdigit()</code> , <code>ispunct()</code> , or <code>isspace()</code> is true. In "C" locale, <code>isalpha()</code> returns true only for the characters for which <code>isupper()</code> or <code>islower()</code> is true.
	<code>isalnum()</code>	Tests for any character for which <code>isalpha()</code> or <code>isdigit()</code> is true (letter or digit).
	<code>isascii()</code>	Tests for any ASCII character, code between 0 and 0177 inclusive.
	<code>isblank()</code>	Tests whether <i>c</i> is a character of class blank in the current locale. This macro/function is not available to applications conforming to standards prior to SUSv3. See <code>standards(5)</code>
	<code>iscntrl()</code>	Tests for any "control character" as defined by the character set.
	<code>isdigit()</code>	Tests for any decimal-digit character.
Default	<code>isgraph()</code>	Tests for any character for which <code>ispunct()</code> , <code>isupper()</code> , <code>islower()</code> , and <code>isdigit()</code> is true.
Standard conforming	<code>isgraph()</code>	Tests for any character for which <code>isalnum()</code> and <code>ispunct()</code> are true, or any character in the current locale-defined "graph" class which is neither a space (" ") nor a character for which <code>iscntrl()</code> is true.
Default	<code>isprint()</code>	Tests for any character for which <code>ispunct()</code> , <code>isupper()</code> , <code>islower()</code> , <code>isdigit()</code> , and the space character (" ") is true.
Standard conforming	<code>isprint()</code>	Tests for any character for which <code>iscntrl()</code> is false, and <code>isalnum()</code> , <code>isgraph()</code> , <code>ispunct()</code> , the space character (" "), and the characters in the current locale-defined "print" class are true.
	<code>ispunct()</code>	Tests for any printing character which is neither a space (" ") nor a character for which <code>isalnum()</code> or <code>iscntrl()</code> is true.
	<code>isspace()</code>	Tests for any space, tab, carriage-return, newline, vertical-tab or form-feed (standard white-space characters) or for one of the current locale-defined set of characters for which <code>isalnum()</code> is false. In the C locale, <code>isspace()</code> returns true only for the standard white-space characters.
	<code>isupper()</code>	Tests for any character that is an upper-case letter or is one of the current locale-defined set of characters for which none of <code>iscntrl()</code> , <code>isdigit()</code> , <code>ispunct()</code> , <code>isspace()</code> , or <code>islower()</code> is true. In the "C" locale, <code>isupper()</code> returns true only for the characters defined as upper-case ASCII characters.
Default	<code>isxdigit()</code>	Tests for any hexadecimal-digit character ([0-9], [A-F], or [a-f]).

ctype(3C)

Standard conforming

`isxdigit()` Tests for any hexadecimal-digit character (`[0-9]`, `[A-F]`, or `[a-f]` or the current locale-defined sets of characters representing the hexadecimal digits 10 to 15 inclusive). In the "C" locale, only
`0 1 2 3 4 5 6 7 8 9 A B C D E F a b c d e f`
are included.

RETURN VALUES

If the argument to any of the character handling macros is not in the domain of the function, the result is undefined. Otherwise, the macro or function returns non-zero if the classification is `TRUE` and 0 if the classification is `FALSE`.

USAGE

These macros or functions can be used safely in multithreaded applications, as long as `setlocale(3C)` is not being called to change the locale.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

SEE ALSO

`setlocale(3C)`, `stdio(3C)`, `ascii(5)`, `environ(5)`, `standards(5)`

NAME	cuserid – get character login name of the user				
SYNOPSIS	<pre>#include <stdio.h> char *cuserid(char *s);</pre>				
DESCRIPTION	<p>The <code>cuserid()</code> function generates a character-string representation of the login name under which the owner of the current process is logged in. If <code>s</code> is a null pointer, this representation is generated in an internal static area whose address is returned. Otherwise, <code>s</code> is assumed to point to an array of at least <code>L_cuserid</code> characters; the representation is left in this array. The constant <code>L_cuserid</code> is defined in the <code><stdio.h></code> header.</p> <p>In multithreaded applications, the caller must always supply an array <code>s</code> for the return value.</p>				
RETURN VALUES	If the login name cannot be found, <code>cuserid()</code> returns a null pointer. If <code>s</code> is not a null pointer, the null character <code>'\0'</code> will be placed at <code>s[0]</code> .				
ATTRIBUTES	See <code>attributes(5)</code> 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>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>getlogin(3C)</code> , <code>getpwnam(3C)</code> , <code>attributes(5)</code>				

dbm(3UCB)

NAME	dbm, dbmopen, dbmclose, fetch, store, delete, firstkey, nextkey – data base subroutines
SYNOPSIS	<pre><code>/usr/ucb/cc [flag ...] file ... -ldb #include <dbm.h> typedef struct { char *dptr; int dsize; } datum; int dbmopen (file); char *file; int dbmclose (); datum fetch (key); datum key; int store (key, dat); datum key, dat; int delete (key); datum key; datum firstkey() datum nextkey (key); datum key;</code></pre>
DESCRIPTION	<p>The <code>dbm()</code> library has been superseded by <code>ndbm</code> (see ndbm(3C)).</p> <p>These functions maintain key/content pairs in a data base. The functions will handle very large (a billion blocks) databases and will access a keyed item in one or two file system accesses.</p> <p><i>key/dat</i> and their content are described by the <code>datum</code> typedef. A <code>datum</code> specifies a string of <i>dsize</i> bytes pointed to by <i>dptr</i>. Arbitrary binary data, as well as normal ASCII strings, are allowed. The data base is stored in two files. One file is a directory containing a bit map and has <code>.dir</code> as its suffix. The second file contains all data and has <code>.pag</code> as its suffix.</p> <p>Before a database can be accessed, it must be opened by <code>dbmopen()</code>. At the time of this call, the files <code>file.dir</code> and <code>file.pag</code> must exist. An empty database is created by creating zero-length <code>.dir</code> and <code>.pag</code> files.</p> <p>A database may be closed by calling <code>dbmclose()</code>. You must close a database before opening a new one.</p>

Once open, the data stored under a key is accessed by `fetch()` and data is placed under a key by `store`. A key (and its associated contents) is deleted by `delete()`. A linear pass through all keys in a database may be made, in an (apparently) random order, by use of `firstkey()` and `nextkey()`. `firstkey()` will return the first key in the database. With any key `nextkey()` will return the next key in the database. This code will traverse the data base:

```
for (key = firstkey; key.dptr != NULL; key = nextkey(key))
```

RETURN VALUES

All functions that return an `int` indicate errors with negative values. A zero return indicates no error. Routines that return a datum indicate errors with a `NULL (0) dptr`.

SEE ALSO

`ar(1)`, `cat(1)`, `cp(1)`, `tar(1)`, `ndbm(3C)`

NOTES

Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.

The `.pag` file will contain holes so that its apparent size may be larger than its actual content. Older versions of the UNIX operating system may create real file blocks for these holes when touched. These files cannot be copied by normal means (`cp(1)`, `cat(1)`, `tar(1)`, `ar(1)`) without filling in the holes.

`dptr` pointers returned by these subroutines point into static storage that is changed by subsequent calls.

The sum of the sizes of a key/content pair must not exceed the internal block size (currently 1024 bytes). Moreover all key/content pairs that hash together must fit on a single block. `store` will return an error in the event that a disk block fills with inseparable data.

`delete()` does not physically reclaim file space, although it does make it available for reuse.

The order of keys presented by `firstkey()` and `nextkey()` depends on a hashing function, not on anything interesting.

There are no interlocks and no reliable cache flushing; thus concurrent updating and reading is risky.

The database files (`file.dir` and `file.pag`) are binary and are architecture-specific (for example, they depend on the architecture's byte order.) These files are not guaranteed to be portable across architectures.

decimal_to_floating(3C)

NAME	decimal_to_floating, decimal_to_single, decimal_to_double, decimal_to_extended, decimal_to_quadruple – convert decimal record to floating-point value				
SYNOPSIS	<pre>#include <floatingpoint.h> void decimal_to_single(single *px, decimal_mode *pm, decimal_record *pd, fp_exception_field_type *ps); void decimal_to_double(double *px, decimal_mode *pm, decimal_record *pd, fp_exception_field_type *ps); void decimal_to_extended(extended *px, decimal_mode *pm, decimal_record *pd, fp_exception_field_type *ps); void decimal_to_quadruple(quadruple *px, decimal_mode *pm, decimal_record *pd, fp_exception_field_type *ps);</pre>				
DESCRIPTION	<p>These functions convert the decimal record <i>pd</i> to a floating-point value <i>px</i> observing the rounding direction specified in <i>pm</i> and setting <i>ps</i> to reflect any floating-point exceptions that occur.</p> <p>When <i>pd->fpclass</i> is <i>fp_zero</i>, <i>fp_infinity</i>, <i>fp_quiet</i>, or <i>fp_signaling</i>, <i>px</i> is set to zero, infinity, a quiet NaN, or a signaling NaN, respectively, with the sign indicated by <i>pd->sign</i>. All other fields in <i>pd</i> are ignored.</p> <p>When <i>pd->fpclass</i> is <i>fp_normal</i> or <i>fp_subnormal</i>, <i>pd->ds</i> must contain a null-terminated string of one or more ASCII digits representing a non-zero integer <i>m</i>, and <i>pd->ndigits</i> must be equal to the length of this string. Then <i>px</i> is set to a correctly rounded approximation to</p> $-1^{**}(\text{pd->sign}) * m * 10^{**}(\text{pd->exponent})$ <p><i>pd->more</i> can be set to a non-zero value to indicate that insignificant trailing digits were omitted from <i>pd->ds</i>. In this case, <i>m</i> is replaced by <i>m + delta</i> in the expression above, where <i>delta</i> is some tiny positive fraction.</p> <p>The converted value is rounded according to the rounding direction specified in <i>pm->rd</i>. <i>pm->df</i> and <i>pm->ndigits</i> are not used.</p> <p>On exit, <i>ps</i> contains a bitwise OR of flags corresponding to any floating-point exceptions that occurred. The only possible exceptions are underflow, overflow, and inexact. If no floating-point exceptions occurred, <i>ps</i> is set to zero.</p>				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<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>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	scanf(3C), string_to_decimal(3C), strtod(3C), attributes(5)				

NAME difftime – computes the difference between two calendar times

SYNOPSIS `#include <time.h>`
`double difftime(time_t time1, time_t time0);`

DESCRIPTION The `difftime()` function computes the difference between two calendar times.

RETURN VALUES The `difftime()` functions returns the difference (*time1-time0*) expressed in seconds as a double.

USAGE The `difftime()` function is provided because there are no general arithmetic properties defined for type `time_t`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [ctime\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

directio(3C)

NAME	directio – provide advice to file system
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/fcntl.h> int directio(int <i>fildes</i>, int <i>advice</i>);</pre>
DESCRIPTION	<p>The <code>directio()</code> function provides advice to the system about the expected behavior of the application when accessing the data in the file associated with the open file descriptor <i>fildes</i>. The system uses this information to help optimize accesses to the file's data. The <code>directio()</code> function has no effect on the semantics of the other operations on the data, though it may affect the performance of other operations.</p> <p>The <i>advice</i> argument is kept per file; the last caller of <code>directio()</code> sets the <i>advice</i> for all applications using the file associated with <i>fildes</i>.</p> <p>Values for <i>advice</i> are defined in <code><sys/fcntl.h></code>.</p> <p>DIRECTIO_OFF Applications get the default system behavior when accessing file data.</p> <p>When an application reads data from a file, the data is first cached in system memory and then copied into the application's buffer (see <code>read(2)</code>). If the system detects that the application is reading sequentially from a file, the system will asynchronously "read ahead" from the file into system memory so the data is immediately available for the next <code>read(2)</code> operation.</p> <p>When an application writes data into a file, the data is first cached in system memory and is written to the device at a later time (see <code>write(2)</code>). When possible, the system increases the performance of <code>write(2)</code> operations by cacheing the data in memory pages. The data is copied into system memory and the <code>write(2)</code> operation returns immediately to the application. The data is later written asynchronously to the device. When possible, the cached data is "clustered" into large chunks and written to the device in a single write operation.</p> <p>The system behavior for <code>DIRECTIO_OFF</code> can change without notice.</p> <p>DIRECTIO_ON The system behaves as though the application is not going to reuse the file data in the near future. In other words, the file data is not cached in the system's memory pages.</p> <p>When possible, data is read or written directly between the application's memory and the device when the data is accessed with <code>read(2)</code> and <code>write(2)</code> operations. When such transfers are not possible, the system switches back to the default behavior, but</p>

just for that operation. In general, the transfer is possible when the application's buffer is aligned on a two-byte (short) boundary, the offset into the file is on a device sector boundary, and the size of the operation is a multiple of device sectors.

This advisory is ignored while the file associated with *fildev* is mapped (see `mmap(2)`).

The system behavior for `DIRECTIO_ON` can change without notice.

RETURN VALUES Upon successful completion, `directio()` returns 0. Otherwise, it returns -1 and sets `errno` to indicate the error.

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

<code>EBADF</code>	The <i>fildev</i> argument is not a valid open file descriptor.
<code>ENOTTY</code>	The <i>fildev</i> argument is not associated with a file system that accepts advisory functions.
<code>EINVAL</code>	The value in <i>advise</i> is invalid.

USAGE Small sequential I/O generally performs best with `DIRECTIO_OFF`.

Large sequential I/O generally performs best with `DIRECTIO_ON`, except when a file is sparse or is being extended and is opened with `O_SYNC` or `O_DSYNC` (see `open(2)`).

The `directio()` function is supported for the NFS and UFS file system types (see `fstyp(1M)`).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `fstyp(1M)`, `mmap(2)`, `open(2)`, `read(2)`, `write(2)`, `fcntl.h(3HEAD)`, `attributes(5)`

WARNINGS Switching between `DIRECTIO_OFF` and `DIRECTIO_ON` can slow the system because each switch to `DIRECTIO_ON` might entail flushing the file's data from the system's memory.

dirname(3C)

NAME | dirname – report the parent directory name of a file path name

SYNOPSIS | #include <libgen.h>

| char *dirname (char *path) ;

DESCRIPTION | The dirname () function takes a pointer to a character string that contains a pathname, and returns a pointer to a string that is a pathname of the parent directory of that file. Trailing '/' characters in the path are not counted as part of the path.

| If path does not contain a '/', then dirname () returns a pointer to the string "." . If path is a null pointer or points to an empty string, dirname () returns a pointer to the string "." .

RETURN VALUES | The dirname () function returns a pointer to a string that is the parent directory of path. If path is a null pointer or points to an empty string, a pointer to a string "." is returned.

ERRORS | No errors are defined.

EXAMPLES | **EXAMPLE 1** Changing the Current Directory to the Parent Directory.

| The following code fragment reads a pathname, changes the current working directory to the parent directory of the named file (see chdir(2)), and opens the file.

| char path[MAXPATHLEN], *pathcopy;

| int fd;

| fgets(path, MAXPATHLEN, stdin);

| pathcopy = strdup(path);

| chdir(dirname(pathcopy));

| fd = open(basename(path), O_RDONLY);

| **EXAMPLE 2** Sample Input and Output Strings for dirname () .

| In the following table, the input string is the value pointed to by path, and the output string is the return value of the dirname () function.

Input String	Output String
"/usr/lib"	"/usr"
"/usr/"	"/"
"usr"	"/"
"/"	"/"
."	."
.."	.."

USAGE | The dirname () function modifies the string pointed to by path.

| The dirname () and basename(3C) functions together yield a complete pathname. The expression dirname (path) obtains the pathname of the directory where basename (path) is found.

dirname(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `basename(1)`, `chdir(2)`, `basename(3C)`, `attributes(5)`, `standards(5)`

div(3C)

NAME	div, ldiv, lldiv – compute the quotient and remainder						
SYNOPSIS	<pre>#include <stdlib.h> div_t div(int numer, int denom); ldiv_t ldiv(long int numer, long int denom); lldiv_t lldiv(long long numer, long long denom);</pre>						
DESCRIPTION	<p>The <code>div()</code> function computes the quotient and remainder of the division of the numerator <i>numer</i> by the denominator <i>denom</i>. It provides a well-defined semantics for the signed integral division and remainder operations, unlike the implementation-defined semantics of the built-in operations. The sign of the resulting quotient is that of the algebraic quotient, and if the division is inexact, the magnitude of the resulting quotient is the largest integer less than the magnitude of the algebraic quotient. If the result cannot be represented, the behavior is undefined; otherwise, <i>quotient * denom + remainder</i> will equal <i>numer</i>.</p> <p>The <code>ldiv()</code> and <code>lldiv()</code> functions are similar to <code>div()</code>, except that the arguments and the members of the returned structure are different. The <code>ldiv()</code> function returns a structure of type <code>ldiv_t</code> and has type <code>long int</code>. The <code>lldiv()</code> function returns a structure of type <code>lldiv_t</code> and has type <code>long long</code>.</p>						
RETURN VALUES	<p>The <code>div()</code> function returns a structure of type <code>div_t</code>, comprising both the quotient and remainder:</p> <pre>int quot; /*quotient*/ int rem; /*remainder*/</pre> <p>The <code>ldiv()</code> function returns a structure of type <code>ldiv_t</code> and <code>lldiv()</code> returns a structure of type <code>lldiv_t</code>, comprising both the quotient and remainder:</p> <pre>long int quot; /*quotient*/ long int rem; /*remainder*/</pre>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>attributes(5)</code> , <code>standards(5)</code>						

NAME	dladdr, dladdr1 – translate address to symbolic information												
SYNOPSIS	<pre>#include <dlfcn.h> int dladdr(void *address, Dl_info *dli); int dladdr1(void *address, Dl_info *dli, void **info, int flags);</pre>												
DESCRIPTION	<p>The <code>dladdr()</code> and <code>dladdr1()</code> functions determine if the specified <i>address</i> is located within one of the mapped objects that make up the current applications address space. An address is deemed to fall within a mapped object when it is between the base address, and the <code>_end</code> address of that object. If a mapped object fits this criteria, the symbol table made available to the runtime linker is searched to locate the nearest symbol to the specified address. The nearest symbol is one that has a value less than or equal to the required address.</p> <p>The <code>Dl_info</code> structure must be preallocated by the user. The structure members are filled in by <code>dladdr()</code> based on the specified <i>address</i>. The <code>Dl_info</code> structure includes the following members:</p> <pre>const char * dli_fname; void * dli_fbase; const char * dli_sname; void * dli_saddr;</pre> <p>Descriptions of these members appear below.</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>dli_fname</code></td> <td>Contains a pointer to the filename of the containing object.</td> </tr> <tr> <td><code>dli_fbase</code></td> <td>Contains the base address of the containing object.</td> </tr> <tr> <td><code>dli_sname</code></td> <td>Contains a pointer to the symbol name nearest to the specified address. This symbol either has the same address, or is the nearest symbol with a lower address.</td> </tr> <tr> <td><code>dli_saddr</code></td> <td>Contains the actual address of the above symbol.</td> </tr> </table> <p>The <code>dladdr1()</code> function provides for addition information to be returned as specified by the <i>flags</i> argument:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>RTLD_DL_SYMENT</code></td> <td>Obtain the ELF symbol table entry for the matched symbol. The <i>info</i> argument points to a symbol pointer as defined in <code><sys/elf.h></code> (<code>Elf32_Sym **info</code> or <code>Elf64_Sym **info</code>).</td> </tr> <tr> <td><code>RTLD_DL_LINKMAP</code></td> <td>Obtain the <code>Link_map</code> for the matched file. The <i>info</i> argument points to a <code>Link_map</code> pointer as defined in <code><sys/link.h></code> (<code>Link_map **info</code>).</td> </tr> </table>	<code>dli_fname</code>	Contains a pointer to the filename of the containing object.	<code>dli_fbase</code>	Contains the base address of the containing object.	<code>dli_sname</code>	Contains a pointer to the symbol name nearest to the specified address. This symbol either has the same address, or is the nearest symbol with a lower address.	<code>dli_saddr</code>	Contains the actual address of the above symbol.	<code>RTLD_DL_SYMENT</code>	Obtain the ELF symbol table entry for the matched symbol. The <i>info</i> argument points to a symbol pointer as defined in <code><sys/elf.h></code> (<code>Elf32_Sym **info</code> or <code>Elf64_Sym **info</code>).	<code>RTLD_DL_LINKMAP</code>	Obtain the <code>Link_map</code> for the matched file. The <i>info</i> argument points to a <code>Link_map</code> pointer as defined in <code><sys/link.h></code> (<code>Link_map **info</code>).
<code>dli_fname</code>	Contains a pointer to the filename of the containing object.												
<code>dli_fbase</code>	Contains the base address of the containing object.												
<code>dli_sname</code>	Contains a pointer to the symbol name nearest to the specified address. This symbol either has the same address, or is the nearest symbol with a lower address.												
<code>dli_saddr</code>	Contains the actual address of the above symbol.												
<code>RTLD_DL_SYMENT</code>	Obtain the ELF symbol table entry for the matched symbol. The <i>info</i> argument points to a symbol pointer as defined in <code><sys/elf.h></code> (<code>Elf32_Sym **info</code> or <code>Elf64_Sym **info</code>).												
<code>RTLD_DL_LINKMAP</code>	Obtain the <code>Link_map</code> for the matched file. The <i>info</i> argument points to a <code>Link_map</code> pointer as defined in <code><sys/link.h></code> (<code>Link_map **info</code>).												
RETURN VALUES	<p>If the specified <i>address</i> cannot be matched to a mapped object, a 0 is returned. Otherwise, a non-zero return is made and the associated <code>Dl_info</code> elements are filled.</p>												

dladdr(3C)

USAGE The `dladdr()` and `dladdr1()` functions are one of a family of functions that give the user direct access to the dynamic linking facilities. These facilities are available to dynamically-linked processes only. See *Linker and Libraries Guide*.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `ld(1)`, `dlclose(3C)`, `dldump(3C)`, `dlerror(3C)`, `dlopen(3C)`, `dlsym(3C)`, `attributes(5)`

Linker and Libraries Guide

NOTES The `Dl_info` pointer elements point to addresses within the mapped objects. These pointers can become invalid if objects are removed prior to these elements use. See `dlclose()`.

If no symbol is found to describe the specified address, both the `dli_sname` and `dli_saddr` members are set to 0.

NAME	dlclose – close a shared object						
SYNOPSIS	<pre>#include <dlfcn.h> int dlclose(void *handle);</pre>						
DESCRIPTION	<p>The <code>dlclose()</code> function decrements the reference count of the supplied <i>handle</i>. This <i>handle</i> represents an executable object file and its dependencies, acquired from a previous call to <code>dlopen()</code>. A <i>handle</i> that is no longer referenced is processed in an attempt to unload any objects that are associated with the <i>handle</i> from the current process. An unreferenced <i>handle</i> is no longer available to <code>dlsym()</code>.</p> <p>Any finalization code within an object is executed prior to that object being unloaded. Any routines registered by an object using <code>atexit(3C)</code> are called prior to that object being unloaded. See NOTES.</p>						
RETURN VALUES	If the <i>handle</i> was successfully unreferenced, <code>dlclose()</code> returns 0. If the <i>handle</i> is invalid, or an error occurred as a result of unloading an object, <code>dlclose()</code> returns a non-zero value. Additional diagnostic information is available through <code>dLError()</code> .						
USAGE	The <code>dlclose()</code> function is one of a family of functions that give the user direct access to the dynamic linking facilities. These facilities are available to dynamically-linked processes only. See the <i>Linker and Libraries Guide</i> .						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<p><code>ld(1)</code>, <code>ld.so.1(1)</code>, <code>atexit(3C)</code>, <code>dladdr(3C)</code>, <code>dldump(3C)</code>, <code>dLError(3C)</code>, <code>dlopen(3C)</code>, <code>dlsym(3C)</code>, <code>attributes(5)</code>, <code>standards(5)</code></p> <p><i>Linker and Libraries Guide</i></p>						
NOTES	<p>A successful invocation of <code>dlclose()</code> does not guarantee that the objects associated with the <i>handle</i> are removed from the address space of the current process. Objects can be referenced by multiple <i>handles</i>, or by other objects. An object is not removed from the address space of the current process until all references to that object are removed.</p> <p>Once an object has been closed by <code>dlclose()</code>, referencing symbols contained in that object can cause undefined behavior.</p>						

dlclose(3C)

As part of unloading an object, finalization code within the object is called *before* the `dlclose()` returns. This finalization is user code, and as such, can produce errors that can not be caught by `dlclose()`. For example, an object loaded using `RTLD_LAZY` that attempts to call a function that can not be located, results in process termination. Erroneous programming practices within the finalization code can also result in process termination. The runtime linker's debugging facility can offer help identifying these types of error. See the `LD_DEBUG` environment variable of `ld.so.1(1)`.

NAME	dldump – create a new file from a dynamic object component of the calling process
SYNOPSIS	<pre>#include <dlfcn.h> int dldump(const char * <i>ipath</i>, const char * <i>opath</i>, int <i>flags</i>);</pre>
DESCRIPTION	<p>The <code>dldump()</code> function creates a new dynamic object <i>opath</i> from an existing dynamic object <i>ipath</i> that is bound to the current process. An <i>ipath</i> value of 0 is interpreted as the dynamic object that started the process. The new object is constructed from the existing objects' disc file. Relocations can be applied to the new object to pre-bind it to other dynamic objects, or fix the object to a specific memory location. In addition, data elements within the new object can be obtained from the objects' memory image as this data exists in the calling process.</p> <p>These techniques allow the new object to be executed with a lower startup cost. This reduction can be because of less relocations being required to load the object, or because of a reduction in the data processing requirements of the object. However, limitations can exist in using these techniques. The application of relocations to the new dynamic object <i>opath</i> can restrict its flexibility within a dynamically changing environment. In addition, limitations in regards to data usage can make dumping a memory image impractical. See EXAMPLES.</p> <p>The runtime linker verifies that the dynamic object <i>ipath</i> is mapped as part of the current process. Thus, the object must either be the dynamic object that started the process, one of the process's dependencies, or an object that has been preloaded. See <code>exec(2)</code>, and <code>ld.so.1(1)</code>.</p> <p>As part of the runtime processing of a dynamic object, <i>relocation</i> records within the object are interpreted and applied to offsets within the object. These offsets are said to be <i>relocated</i>. Relocations can be categorized into two basic types: <i>non-symbolic</i> and <i>symbolic</i>.</p> <p>The <i>non-symbolic</i> relocation is a simple <i>relative</i> relocation that requires the base address at which the object is mapped to perform the relocation. The <i>symbolic</i> relocation requires the address of an associated symbol, and results in a <i>binding</i> to the dynamic object that defines this symbol. The symbol definition can originate from any of the dynamic objects that make up the process, that is, the object that started the process, one of the process's dependencies, an object that has been preloaded, or the dynamic object being relocated.</p> <p>The <i>flags</i> parameter controls the relocation processing and other attributes of producing the new dynamic object <i>opath</i>. Without any <i>flags</i>, the new object is constructed solely from the contents of the <i>ipath</i> disc file without any relocations applied.</p> <p>Various relocation flags can be <i>or'</i>ed into the <i>flags</i> parameter to affect the relocations that are applied to the new object. <i>Non-symbolic</i> relocations can be applied using the following:</p> <pre>RTLD_REL_RELATIVE</pre> <p>Relocation records from the object <i>ipath</i>, that define <i>relative</i> relocations, are applied to the object <i>opath</i>.</p>

dldump(3C)

A variety of *symbolic* relocations can be applied using the following flags (each of these flags also implies `RTLDR_REL_RELATIVE` is in effect):

`RTLDR_REL_EXEC`

Symbolic relocations that result in binding *ipath* to the dynamic object that started the process, commonly a dynamic executable, are applied to the object *opath*.

`RTLDR_REL_DEPENDS`

Symbolic relocations that result in binding *ipath* to any of the dynamic dependencies of the process are applied to the object *opath*.

`RTLDR_REL_PRELOAD`

Symbolic relocations that result in binding *ipath* to any objects preloaded with the process are applied to the object *opath*. See `LD_PRELOAD` in `ld.so.1(1)`.

`RTLDR_REL_SELF`

Symbolic relocations that result in binding *ipath* to itself, are applied to the object *opath*.

`RTLDR_REL_WEAK`

Weak relocations that remain unresolved are applied to the object *opath* as 0.

`RTLDR_REL_ALL`

All relocation records defined in the object *ipath* are applied to the new object *opath*. This is basically a concatenation of all the above relocation flags.

Note that for dynamic executables, `RTLDR_REL_RELATIVE`, `RTLDR_REL_EXEC`, and `RTLDR_REL_SELF` have no effect. See `EXAMPLES`.

If relocations, knowledgeable of the base address of the mapped object, are applied to the new object *opath*, then the new object becomes fixed to the location that the *ipath* image is mapped within the current process.

Any relocations applied to the new object *opath* will have the original relocation record removed so that the relocation will not be applied more than once. Otherwise, the new object *opath* will retain the relocation records as they exist in the *ipath* disc file.

The following additional attributes for creating the new dynamic object *opath* can be specified using the *flags* parameter:

`RTLDR_MEMORY`

The new object *opath* is constructed from the current memory contents of the *ipath* image as it exists in the calling process. This option allows data modified by the calling process to be captured in the new object. Note that not all data modifications may be applicable for capture; significant restrictions exist in using this technique. See `EXAMPLES`. By default, when processing a dynamic executable, any allocated memory that follows the end of the data segment is captured in the new object (see `malloc(3C)` and `brk(2)`). This data, which represents the process heap, is saved as a new `.SUNW_heap` section in the object *opath*. The objects' program headers and symbol entries, such as `_end`, are adjusted accordingly. See also `RTLDR_NOHEAP`.

When using this attribute, any relocations that have been applied to the *ipath* memory image that do not fall into one of the requested relocation categories are undone, that is, the relocated element is returned to the value as it existed in the *ipath* disc file.

RTLD_STRIP

Only collect allocatable sections within the object *opath*. Sections that are not part of the dynamic objects' memory image are removed. **RTLD_STRIP** reduces the size of the *opath* disc file and is comparable to having run the new object through `strip(1)`.

RTLD_NOHEAP

Do not save any heap to the new object. This option is only meaningful when processing a dynamic executable with the **RTLD_MEMORY** attribute and allows for reducing the size of the *opath* disc file. The executable must confine its data initialization to data elements within its data segment, and must not use any allocated data elements that comprise the heap.

It should be emphasized, that an object created by `dldump()` is simply an updated ELF object file. No additional state regarding the process at the time `dldump()` is called is maintained in the new object. `dldump()` does not provide a panacea for checkpoint and resume. A new dynamic executable, for example, will not start where the original executable called `dldump()`. It will gain control at the executable's normal entry point. See **EXAMPLES**.

RETURN VALUES

On successful creation of the new object, `dldump()` returns 0. Otherwise, a non-zero value is returned and more detailed diagnostic information is available through `dlerror()`.

EXAMPLES

EXAMPLE 1 Sample code using `dldump()`.

The following code fragment, which can be part of a dynamic executable `a.out`, can be used to create a new shared object from one of the dynamic executables' dependencies `libfoo.so.1`:

```
const char *   ipath = "libfoo.so.1";
const char *   opath = "./tmp/libfoo.so.1";
...
if (dldump(ipath, opath, RTLD_REL_RELATIVE) != 0)
    (void) printf("dldump failed: %s\n", dlerror());
```

The new shared object *opath* is fixed to the address of the mapped *ipath* bound to the dynamic executable `a.out`. All relative relocations are applied to this new shared object, which will reduce its relocation overhead when it is used as part of another process.

By performing only relative relocations, any symbolic relocation records remain defined within the new object, and thus the dynamic binding to external symbols will be preserved when the new object is used.

EXAMPLE 1 Sample code using `dldump()`. (Continued)

Use of the other relocation flags can fix specific relocations in the new object and thus can reduce even more the runtime relocation startup cost of the new object. However, this will also restrict the flexibility of using the new object within a dynamically changing environment, as it will bind the new object to some or all of the dynamic objects presently mapped as part of the process.

For example, the use of `RTLD_REL_SELF` will cause any references to symbols from *ipath* to be bound to definitions within itself if no other preceding object defined the same symbol. In other words, a call to *foo()* within *ipath* will bind to the definition *foo* within the same object. Therefore, *opath* will have one less binding that must be computed at runtime. This reduces the startup cost of using *opath* by other applications; however, interposition of the symbol *foo* will no longer be possible.

Using a dumped shared object with applied relocations as an applications dependency normally requires that the application have the same dependencies as the application that produced the dumped image. Dumping shared objects, and the various flags associated with relocation processing, have some specialized uses. However, the technique is intended as a building block for future technology.

The following code fragment, which is part of the dynamic executable `a.out`, can be used to create a new version of the dynamic executable:

```
static char *      dumped = 0;
const char *      opath = "./a.out.new";
...
if (dumped == 0) {
    char          buffer[100];
    int           size;
    time_t        seconds;
    ...
    /* Perform data initialization */
    seconds = time((time_t *)0);
    size = cftime(buffer, (char *)0, &seconds);
    if ((dumped = (char *)malloc(size + 1)) == 0) {
        (void) printf("malloc failed: %s\n", strerror(errno));
        return (1);
    }
    (void) strcpy(dumped, buffer);
    ...
    /*
     * Tear down any undesirable data initializations and
     * dump the dynamic executables memory image.
     */
    _exithandle( );
    _exit(dldump(0, opath, RTLD_MEMORY));
}
(void) printf("Dumped: %s\n", dumped);
```


EXAMPLE 1 Sample code using `dldump()`. (Continued)

Any modifications made to the dynamic executable, up to the point the `dldump()` call is made, are saved in the new object `a.out.new`. This mechanism allows the executable to update parts of its data segment and heap prior to creating the new object. In this case, the date the executable is dumped is saved in the new object. The new object can then be executed without having to carry out the same (presumably expensive) initialization.

For greatest flexibility, this example does not save *any* relocated information. The elements of the dynamic executable `ipath` that have been modified by relocations at process startup, that is, references to external functions, are returned to the values of these elements as they existed in the `ipath` disc file. This preservation of relocation records allows the new dynamic executable to be flexible, and correctly bind and initialize to its dependencies when executed on the same or newer upgrades of the OS.

Fixing relocations by applying some of the relocation flags would bind the new object to the dependencies presently mapped as part of the process calling `dldump()`. It may also remove necessary copy relocation processing required for the correct initialization of its shared object dependencies. Therefore, if the new dynamic executables' dependencies have no specialized initialization requirements, the executable may still only interact correctly with the dependencies to which it binds if they were mapped to the same locations as they were when `dldump()` was called.

Note that for dynamic executables, `RTLD_REL_RELATIVE`, `RTLD_REL_EXEC`, and `RTLD_REL_SELF` have no effect, as relocations within the dynamic executable will have been fixed when it was created by `ld(1)`.

When `RTLD_MEMORY` is used, care should be taken to insure that dumped data sections that reference external objects are not reused without appropriate re-initialization. For example, if a data item contains a file descriptor, a variable returned from a shared object, or some other external data, and this data item has been initialized prior to the `dldump()` call, its value will have no meaning in the new dumped image.

When `RTLD_MEMORY` is used, any modification to a data item that is initialized via a relocation whose relocation record will be retained in the new image will effectively be lost or invalidated within the new image. For example, if a pointer to an external object is incremented prior to the `dldump()` call, this data item will be reset to its disc file contents so that it can be relocated when the new image is used; hence, the previous increment is lost.

Non-idempotent data initializations may prevent the use of `RTLD_MEMORY`. For example, the addition of elements to a linked-list via `init` sections can result in the linked-list data being captured in the new image. Running this new image may result in `init` sections continuing to add new elements to the list without the prerequisite initialization of the list head. It is recommended that `_exithandle(3C)` be called before `dldump()` to tear down any data initializations established via initialization code. Note that this may invalidate the calling image; thus, following the call to

dldump(3C)

EXAMPLE 1 Sample code using `dldump()`. (Continued)

`dldump()`, only a call to `_exit(2)` should be made.

USAGE The `dldump()` function is one of a family of functions that give the user direct access to the dynamic linking facilities. These facilities are available to dynamically-linked processes only. See *Linker and Libraries Guide*.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcsu
MT-Level	MT-Safe

SEE ALSO `ld(1)`, `ld.so.1(1)`, `strip(1)`, `_exit(2)`, `brk(2)`, `exec(2)`, `_exithandle(3C)`, `dladdr(3C)`, `dlclose(3C)`, `dLError(3C)`, `dlopen(3C)`, `dlsym(3C)`, `end(3C)`, `malloc(3C)`, `attributes(5)`

Linker and Libraries Guide

NOTES These functions are available to dynamically-linked processes only.

Any NOBITS sections within the *ipath* are expanded to PROGBITS sections within the *opath*. NOBITS sections occupy no space within an ELF file image. NOBITS sections declare memory that must be created and zero-filled when the object is mapped into the runtime environment. `.bss` is a typical example of this section type. PROGBITS sections, on the other hand, hold information defined by the object within the ELF file image. This section conversion reduces the runtime initialization cost of the new dumped object but increases the objects' disc space requirement.

When a shared object is dumped, and relocations are applied which are knowledgeable of the base address of the mapped object, the new object is fixed to this new base address. The dumped object has its ELF type reclassified to be a dynamic executable. The dumped object can be processed by the runtime linker, but is not valid as input to the link-editor.

If relocations are applied to the new object, any remaining relocation records are reorganized for better locality of reference. The relocation sections are renamed to `.SUNW_reloc` and the association with the section to relocate, is lost. Only the offset of the relocation record is meaningful. `.SUNW_reloc` relocations do not make the new object invalid to either the runtime linker or link-editor, but can reduce the objects analysis with some ELF readers.

NAME	dlerror – get diagnostic information						
SYNOPSIS	<pre>#include <dlfcn.h> char *dlerror(void);</pre>						
DESCRIPTION	The <code>dlerror()</code> function returns a null-terminated character string that describes the last error that occurred during dynamic linking processing. The returned string contains no trailing newline. If no dynamic linking errors have occurred since the last invocation of <code>dlerror()</code> , <code>dlerror()</code> returns NULL. Thus, invoking <code>dlerror()</code> a second time, immediately following a prior invocation, results in NULL being returned.						
USAGE	The <code>dlerror()</code> function is one of a family of functions that give the user direct access to the dynamic linking facilities. These facilities are available to dynamically-linked processes only. See <i>Linker and Libraries Guide</i> .						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>ld(1)</code> , <code>dladdr(3C)</code> , <code>dlclose(3C)</code> , <code>dl_dump(3C)</code> , <code>dlopen(3C)</code> , <code>dlsym(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code> <i>Linker and Libraries Guide</i>						
NOTES	The messages returned by <code>dlerror()</code> can reside in a static buffer that is overwritten on each call to <code>dlerror()</code> . Application code should not write to this buffer. Programs wanting to preserve an error message should make their own copies of that message.						

dldlfo(3C)

NAME	dldlfo – dynamic load information
SYNOPSIS	<pre>#include <dldfcn.h> #include <link.h> #include <limits.h> int dldlfo(void *handle, int request, void *p);</pre>
DESCRIPTION	<p>The dldlfo() function sets or extracts information from the runtime linker ld.so.1(1). This function is loosely modeled after the ioctl(2) function. The <i>request</i> argument and a third argument of varying type are passed to dldlfo(). The action taken by dldlfo() depends on the value of the <i>request</i> that is provided.</p> <p>The <i>handle</i> argument is either the value that is returned from a dlopen(3C) or dllopen() call, or the special handle RTLD_SELF. A <i>handle</i> argument is required for all requests except RTLD_DI_CONFIGADDR, RTLD_DI_GETSIGNAL, and RTLD_DI_SETSIGNAL. If <i>handle</i> is the value returned from a dlopen() or dllopen() call, the information returned by the dldlfo() call pertains to the specified object. If <i>handle</i> is the special handle RTLD_SELF, the information returned by the dldlfo() call pertains to the caller.</p> <p>The <i>request</i> argument can take the following values:</p> <p>RTLD_DI_CONFIGADDR Obtain the configuration file information. The <i>p</i> argument is a Dl_info pointer (Dl_info *p). The following elements from this structure are initialized:</p> <ul style="list-style-type: none">dli_fname The full name of the configuration file.dli_fbase The base address of the configuration file loaded into memory. <p>RTLD_DI_LINKMAP Obtain the Link_map for the <i>handle</i> that is specified. The <i>p</i> argument points to a Link_map pointer (Link_map **p). The actual storage for the Link_map structure is maintained by ld.so.1.</p> <p>The Link_map structure includes the following members:</p> <pre>unsigned long l_addr; /* base address */ char *l_name; /* object name */ Elf32_Dyn *l_ld; /* .dynamic section */ Link_map *l_next; /* next link object */ Link_map *l_prev; /* previous link object */ char *l_refname; /* filter reference name */</pre> <ul style="list-style-type: none">l_addr The base address of the object loaded into memory.l_name The full name of the loaded object. This full name is the filename of the object as referenced by ld.so.1.l_ld Points to the SHT_DYNAMIC structure.

l_next

The next `Link_map` on the link-map list. Other objects on the same link-map list as the current object can be examined by following the `l_next` and `l_prev` members.

l_prev

The previous `Link_map` on the link-map list.

l_refname

If the object that is referenced is a *filter*, this member points to the name of the object being filtered. If the object is not a *filter*, this member is 0. See the *Linker and Libraries Guide*.

RTLD_DI_LMID

Obtain the ID for the link-map list upon which the *handle* is loaded. The *p* argument is a `Lmid_t` pointer (`Lmid_t *p`).

RTLD_DI_SERINFO

Obtain the library search paths for the *handle* that is specified. The *p* argument is a `Dl_serinfo` pointer (`Dl_serinfo *p`). A user must first initialize the `Dl_serinfo` structure with a `RTLD_DI_SERINFO` request. See `EXAMPLES`.

The returned `Dl_serinfo` structure contains `dls_cnt` `Dl_serpath` entries. Each entry's `dlp_name` member points to the search path. The corresponding `dlp_info` member contains one or more flags indicating the origin of the path. See the `LA_SER_*` flags that are defined in `<link.h>`.

RTLD_DI_SERINFO

Initialize a `Dl_serinfo` structure for use in a `RTLD_DI_SERINFO` request. Both the `dls_cnt` and `dls_size` members are returned. The `dls_cnt` member indicates the number of search paths that are applicable to the *handle*. The `dls_size` member indicates the total size of a `Dl_serinfo` buffer required to hold `dls_cnt` `Dl_serpath` entries and the associated search path strings.

To obtain the complete path information, a new `Dl_serinfo` buffer of size `dls_size` should be allocated. This new buffer should be initialized with the `dls_cnt` and `dls_size` entries and passed to a `RTLD_DI_SERINFO` request. See `EXAMPLES`.

RTLD_DI_ORIGIN

Obtain the origin of the dynamic object that is associated with the *handle*. The *p* argument is a `char` pointer (`char *p`). The `dirname(3C)` of the associated object's `realpath(3C)`, which can be no larger than `{PATH_MAX}`, is copied to the pointer *p*.

RTLD_DI_GETSIGNAL

Obtain the numeric signal number used by the runtime linker to kill the process in the event of a fatal runtime error. The *p* argument is an `int` pointer (`int *p`). The signal number is copied to the pointer *p*.

By default, the signal used by the runtime linker to terminate a process is `SIGKILL`. See `thr_kill(3C)`. This default can be changed by calling `dldlfo()` with `RTLD_DI_SETSIGNAL` or by setting the environment variable `LD_SIGNAL`. See `ld.so.1(1)`.

dldlfo(3C)

RTLD_DI_SETSIGNAL

Provide a numeric signal number used by the runtime linker to kill the process in the event of a fatal runtime error. The *p* argument is an int pointer (int *p). The value pointed to by *p* is established as the terminating signal value.

The current signal number used by the runtime linker to terminate a process can be obtained from dldlfo() using RTLD_DI_GETSIGNAL. Use of the RTLD_DI_SETSIGNAL option is equivalent to setting the environment variable LD_SIGNAL. See ld.so.1(1).

RETURN VALUES

The dldlfo() function returns -1 if the *request* is invalid, the parameter *p* is NULL, or the Dl_serinfo structure is uninitialized for a RTLD_DI_SERINFO request. dldlfo() also returns -1 if the *handle* argument does not refer to a valid object opened by dlopen(), or is not the special handle RTLD_SELF. Detailed diagnostic information is available with dldlerror(3C).

EXAMPLES

EXAMPLE 1 Use dldlfo() to obtain the library search paths.

The following example demonstrates how a dynamic object can inspect the library search paths that would be used to locate a simple filename with dlopen(). For simplicity, error checking has been omitted.

```
Dl_serinfo    _info, *info = &_info;
Dl_serpath    *path;
uint_t        cnt;

/* determine search path count and required buffer size */
dldlfo(RTLD_SELF, RTLD_DI_SERINFO, (void *)info);

/* allocate new buffer and initialize */
info = malloc(_info.dls_size);
info->dls_size = _info.dls_size;
info->dls_cnt = _info.dls_cnt;

/* obtain search path information */
dldlfo(RTLD_SELF, RTLD_DI_SERINFO, (void *)info);

path = &info->dls_serpath[0];

for (cnt = 1; cnt <= info->dls_cnt; cnt++, path++) {
    (void) printf("%2d: %s\n", cnt, path->dls_name);
}
```

USAGE

The dldlfo() function is one of a family of functions that give the user direct access to the dynamic linking facilities. These facilities are available to dynamically-linked processes only. See the *Linker and Libraries Guide*.

ATTRIBUTES

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

dldump(3C)

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

SEE ALSO ld(1), ld.so.1(1), ioctl(2), dirname(3C), dlclose(3C), dldump(3C), dlerror(3C), dlopen(3C), dlsym(3C), realpath(3C), thr_kill(3C), attributes(5)

Linker and Libraries Guide

dlopen(3C)

NAME	dlopen, dlmopen – gain access to an executable object file		
SYNOPSIS	<pre>#include <dlfcn.h> #include <link.h> void * dlopen(const char *<i>pathname</i>, int <i>mode</i>); void * dlmopen(Lmid_t <i>lmid</i>, const char *<i>pathname</i>, int <i>mode</i>);</pre>		
DESCRIPTION	<p>The <code>dlopen()</code> function makes an executable object file available to a running process. <code>dlopen()</code> returns to the process a <i>handle</i> that the process can use on subsequent calls to <code>dlsym(3C)</code>, <code>dladdr(3C)</code>, <code>dlinfo(3C)</code>, and <code>dlclose(3C)</code>. The value of this <i>handle</i> should not be interpreted in any way by the process. The <i>pathname</i> argument is the path name of the object to be opened. A path name containing an embedded '/' is interpreted as an absolute path or relative to the current directory. Otherwise, the set of search paths currently in effect by the runtime linker are used to locate the specified file. See NOTES.</p> <p>The <code>dlopen()</code> function also loads any dependencies recorded within <i>pathname</i>. These dependencies are searched in the order in which the dependencies were loaded to locate any additional dependencies. This process continues until all the dependencies of <i>pathname</i> are loaded. This dependency tree is referred to as a group.</p> <p>If the value of <i>pathname</i> is 0, <code>dlopen()</code> provides a <i>handle</i> on a set of global symbol objects. These objects consist of the original program image file, any dependencies loaded at program startup, and any objects loaded using <code>dlopen()</code> with the <code>RTLD_GLOBAL</code> flag. Because the latter set of objects can change during process execution, the set identified by <i>handle</i> can also change dynamically.</p> <p>The <i>mode</i> argument describes how <code>dlopen()</code> operates on <i>pathname</i> with respect to the processing of reference relocations. The <i>mode</i> also affects the scope of visibility of the symbols provided by <i>pathname</i> and its dependencies. This visibility can affect how the resulting <i>handle</i> is used.</p> <p>When an object is loaded, the object can contain references to symbols whose addresses are not known until the object is loaded. These references must be relocated before the symbols can be accessed. References are categorized as either <i>immediate</i> or <i>lazy</i>. Immediate references are typically references to data items used by the object code. Immediate references include pointers to functions and calls to functions made from position-dependent shared objects. Lazy references are typically calls to global functions that are made from position-independent shared objects. The <i>mode</i> argument governs when these references take place. The <i>mode</i> argument can be one of the following values:</p> <table><tr><td><code>RTLD_LAZY</code></td><td>Only immediate symbol references are relocated when the object is first loaded. Lazy references are not relocated until a given function is called for the first time. This value for <i>mode</i> should improve performance, since a process might not require all lazy references in any given object. This behavior mimics the normal loading of dependencies during process initialization. See NOTES.</td></tr></table>	<code>RTLD_LAZY</code>	Only immediate symbol references are relocated when the object is first loaded. Lazy references are not relocated until a given function is called for the first time. This value for <i>mode</i> should improve performance, since a process might not require all lazy references in any given object. This behavior mimics the normal loading of dependencies during process initialization. See NOTES.
<code>RTLD_LAZY</code>	Only immediate symbol references are relocated when the object is first loaded. Lazy references are not relocated until a given function is called for the first time. This value for <i>mode</i> should improve performance, since a process might not require all lazy references in any given object. This behavior mimics the normal loading of dependencies during process initialization. See NOTES.		

RTLD_NOW All necessary relocations are performed when the object is first loaded. This process might waste some processing if relocations are performed for lazy references that are never used. However, this mode ensures that when an object is loaded, all symbols referenced during execution are available. This behavior mimics the loading of dependencies when the environment variable `LD_BIND_NOW` is in effect.

See the *Linker and Libraries Guide* for more information about symbol references.

The visibility of symbols that are available for relocation can be affected by *mode*. To specify the scope of visibility for symbols that are loaded with a `dlopen()` call, *mode* should be a bitwise-inclusive OR with one of the following values:

RTLD_GLOBAL The object's global symbols are made available for the relocation processing of any other object. In addition, symbol lookup using `dlopen(0, mode)` and an associated `dlsym()` allows objects that are loaded with `RTLD_GLOBAL` to be searched.

RTLD_LOCAL The object's global symbols are only available for the relocation processing of other objects that include the same group.

The program image file and any objects loaded at program startup have the mode `RTLD_GLOBAL`. The mode `RTLD_LOCAL` is the default mode for any objects that are acquired with `dlopen()`. A local object can be a dependency of more than one group. Any object of mode `RTLD_LOCAL` that is referenced as a dependency of an object of mode `RTLD_GLOBAL` is promoted to `RTLD_GLOBAL`. In other words, the `RTLD_LOCAL` mode is ignored.

Any object loaded by `dlopen()` that requires relocations against global symbols can reference the symbols in any `RTLD_GLOBAL` object. Objects of this mode are at least the program image file and any objects loaded at program startup. A loaded object can also reference symbols from itself, and from any dependencies the object references. However, the *mode* parameter can also be a bitwise-inclusive OR with one of the following values to affect the scope of symbol availability:

RTLD_GROUP Only symbols from the associated group are made available for relocation. A group is established from the defined object and all the dependencies of that object. A group must be completely self-contained. All dependency relationships between the members of the group must be sufficient to satisfy the relocation requirements of each object that defines the group.

RTLD_PARENT The symbols of the object initiating the `dlopen()` call are made available to the objects obtained by `dlopen()`. This option is useful when hierarchical `dlopen()` families are created. Although the parent object can supply symbols for the relocation of this object, the parent object is not available to `dlsym()` through the returned *handle*.

dlopen(3C)

`RTLD_WORLD` Only symbols from `RTLD_GLOBAL` objects are made available for relocation.

The default modes for `dlopen()` are both `RTLD_WORLD` and `RTLD_GROUP`. If an object requires additional modes, the *mode* parameter can be the bitwise-inclusive OR of the required modes together with the default modes.

The following modes provide additional capabilities outside of relocation processing:

`RTLD_NODELETE` The specified object is tagged to prevent its deletion from the address space as part of a `dlclose()`.

`RTLD_NOLOAD` The specified object is not loaded as part of the `dlopen()`. However, a valid *handle* is returned if the object already exists as part of the process address space. Additional modes can be specified as a bitwise-inclusive OR with the present mode of the object and its dependencies. The `RTLD_NOLOAD` mode provides a means of querying the presence or promoting the modes of an existing dependency.

The default use of a *handle* with `dlsym()` allows a symbol search to inspect all objects that are associated with the group of objects that are loaded from `dlopen()`. The *mode* parameter can also be a bitwise-inclusive OR with the following value to restrict this symbol search:

`RTLD_FIRST` Use of this *handle* with `dlsym()`, restricts the symbol search to the first object associated with the *handle*.

An object can be accessed from a process both with and without `RTLD_FIRST`. Although the object will only be loaded once, two different *handles* are created to provide for the different `dlsym()` requirements.

The `dlmopen()` function is identical to `dlopen()`, except that an identifying link-map ID (*lmid*) is provided. This link-map ID informs the dynamic linking facilities upon which link-map list to load the object. See the *Linker and Libraries Guide* for details about link-maps.

The *lmid* passed to `dlmopen()` identifies the link-map list on which the object is loaded. This parameter can be any valid `Lmid_t` returned by `dlinfo()` or one of the following special values:

`LM_ID_BASE` Load the object on the applications link-map list.

`LM_ID_LDSO` Load the object on the dynamic linkers (`ld.so.1`) link-map list.

`LM_ID_NEWLM` Cause the object to create a new link-map list as part of loading. Objects that are opened on a new link-map list must express all of their dependencies.

RETURN VALUES The `dlopen()` function returns `NULL` if *pathname* cannot be found, cannot be opened for reading, or is not a shared object or a relocatable object. `dlopen()` also returns `NULL` if an error occurs during the process of loading *pathname* or relocating its symbolic references. See **NOTES**. Additional diagnostic information is available through `dlerror()`.

USAGE The `dlopen()` and `dlopen()` functions are members of a family of functions that give the user direct access to the dynamic linking facilities. This family of functions is available only to dynamically-linked processes. See the *Linker and Libraries Guide*.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `ld(1)`, `ld.so.1(1)`, `dladdr(3C)`, `dlclose(3C)`, `dlDump(3C)`, `dlerror(3C)`, `dlinfo(3C)`, `dlsym(3C)`, `attributes(5)`, `standards(5)`

Linker and Libraries Guide

NOTES If *pathname* has dependencies on other objects, these objects are automatically loaded by `dlopen()`. The directory search path used to find *pathname* and any dependencies can be affected by setting the environment variable `LD_LIBRARY_PATH`. Any `LD_LIBRARY_PATH` variable is analyzed once at process startup. The search path can also be affected from a `runpath` setting within the object from which the call to `dlopen()` originates. These search rules will only be applied to path names that do not contain an embedded `'/'`. Objects whose names resolve to the same absolute path name or relative path name can be opened any number of times using `dlopen()`. However, the object that is referenced will only be loaded once into the address space of the current process.

When loading shared objects, the application should open a specific version of the shared object. Do not rely on the version of the shared object pointed to by the symbolic link.

When building objects to be loaded on a new link-map list, some precautions need to be taken. In general, all dependencies must be included when building an object. Also, include `/usr/lib/libmapmalloc.so.1` before `/lib/libc.so.1` when building an object.

When an object is loaded on a new link-map list, the object is isolated from the main running program. Certain global resources are only usable from one link-map list. A few examples are the `sbrk()` based `malloc()`, `libthread()`, and the signal vectors. Care must be taken not to use any of these resources other than from the primary link-map list. These issues are discussed in further detail in the *Linker and Libraries Guide*.

dlopen(3C)

Some symbols defined in dynamic executables or shared objects can not be available to the runtime linker. The symbol table created by `ld` for use by the runtime linker might contain only a subset of the symbols that are defined in the object.

As part of loading a new object, initialization code within the object is called *before* the `dlopen()` returns. This initialization is user code, and as such, can produce errors that can not be caught by `dlopen()`. For example, an object loaded using `RTLD_LAZY` that attempts to call a function that can not be located results in process termination. Erroneous programming practices within the initialization code can also result in process termination. The runtime linker's debugging facility can offer help identifying these types of error. See the `LD_DEBUG` environment variable of `ld.so.1(1)`.

NAME	dlsym – get the address of a symbol in a shared object or executable
SYNOPSIS	<pre>#include <dlfcn.h> void *dlsym(void *restrict handle, const char *restrict name);</pre>
DESCRIPTION	<p>The <code>dlsym()</code> function allows a process to obtain the address of a symbol that is defined within a shared object or executable. The <i>handle</i> argument is either the value returned from a call to <code>dlopen()</code> or one of a family of special handles. The <i>name</i> argument is the symbol's name as a character string.</p> <p>If <i>handle</i> is returned from <code>dlopen()</code>, the associated shared object must not have been closed using <code>dlclose()</code>. A <i>handle</i> can be obtained from <code>dlopen()</code> using the <code>RTLD_FIRST</code> mode. With this mode, the <code>dlsym()</code> function searches for the named symbol in the initial object referenced by <i>handle</i>. Without this mode, the <code>dlsym()</code> function searches for the named symbol in the group of shared objects loaded automatically as a result of loading the object referenced by <i>handle</i>. See dlopen(3C).</p> <p>The following special handles are supported.</p> <p>RTLD_DEFAULT</p> <p>Instructs <code>dlsym()</code> to search for the named symbol starting with the first object loaded, typically the dynamic executable. The search continues through the list of initial dependencies that are loaded with the process, followed by any objects obtained with dlopen(3C). This search follows the default model that is used to relocate all objects within the process.</p> <p>This model also provides for transitioning into a lazy loading environment. If a symbol can not be found in the presently loaded objects, any pending lazy loaded objects are processed in an attempt to locate the symbol. This loading compensates for objects that have not fully defined their dependencies. However, this compensation can undermine the advantages of lazy loading.</p> <p>RTLD_PROBE</p> <p>Instructs <code>dlsym()</code> to search for the named symbol in the same manner as occurs with a <i>handle</i> of <code>RTLD_DEFAULT</code>. However, this model only searches for symbols in the presently loaded objects, together with any lazy loadable objects specifically identified by the caller to provide the named symbol. This handle does not trigger an exhaustive load of any lazy loadable symbols in an attempt to find the named symbol. This handle can provide a more optimal search than would occur using <code>RTLD_DEFAULT</code>.</p> <p>RTLD_NEXT</p> <p>Instructs <code>dlsym()</code> to search for the named symbol in the objects that were loaded following the object from which the <code>dlsym()</code> call is being made.</p> <p>RTLD_SELF</p> <p>Instructs <code>dlsym()</code> to search for the named symbol in the objects that were loaded starting with the object from which the <code>dlsym()</code> call is being made.</p> <p>When used with a special handle, <code>dlsym()</code> is selective in searching objects that have been loaded using <code>dlopen()</code>. These objects are searched for symbols if one of the following conditions are true.</p>

dlsym(3C)

- The object is part of the same local `dlopen()` dependency hierarchy as the calling object. See the *Linker and Libraries Guide* for a description of `dlopen()` dependency hierarchies.
- The object has global search access. See [dlopen\(3C\)](#) for a discussion of the `RTLD_GLOBAL` mode.

RETURN VALUES The `dlsym()` function returns `NULL` if *handle* does not refer to a valid object opened by `dlopen()`, or is not one of the special handles. The function also returns `NULL` if the named symbol cannot be found within any of the objects associated with *handle*. Additional diagnostic information is available through [dlerror\(3C\)](#).

EXAMPLES **EXAMPLE 1** Use `dlopen()` and `dlsym()` to access a function or data objects.

The following code fragment demonstrates how to use `dlopen()` and `dlsym()` to access either function or data objects. For simplicity, error checking has been omitted.

```
void      *handle;
int       *iptr, (*fptr)(int);

/* open the needed object */
handle = dlopen("/usr/home/me/libfoo.so.1", RTLD_LAZY);

/* find the address of function and data objects */
fptr = (int (*)(int))dlsym(handle, "my_function");
iptr = (int *)dlsym(handle, "my_object");

/* invoke function, passing value of integer as a parameter */
(*fptr)(*iptr);
```

EXAMPLE 2 Use `dlsym()` to verify that a particular function is defined.

The following code fragment shows how to use `dlsym()` to verify that a function is defined. If the function exists, the function is called.

```
int (*fptr)();

if ((fptr = (int (*)())dlsym(RTLD_DEFAULT,
    "my_function")) != NULL) {
    (*fptr)();
}
```

USAGE The `dlsym()` function is one of a family of functions that give the user direct access to the dynamic linking facilities. These facilities are available to dynamically-linked processes only. See the *Linker and Libraries Guide*.

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

dlsym(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO ld(1), ld.so.1(1), dladdr(3C), dlclose(3C), dlclose(3C), dlerror(3C),
dlopen(3C), dlopen(3C), attributes(5), standards(5)

Linker and Libraries Guide

drand48(3C)

NAME	drand48, erand48, lrand48, nrand48, mrand48, jrand48, srand48, seed48, lcong48 – generate uniformly distributed pseudo-random numbers
SYNOPSIS	<pre>#include <stdlib.h> double drand48(void); double erand48(unsigned short <i>x_i[3]</i>); long lrnd48(void); long nrand48(unsigned short <i>x_i[3]</i>); long mrnd48(void); long jrnd48(unsigned short <i>x_i[3]</i>); void srand48(long <i>seedval</i>); unsigned short *seed48(unsigned short <i>seed16v[3]</i>); void lcong48(unsigned short <i>param[7]</i>);</pre>
DESCRIPTION	<p>This family of functions generates pseudo-random numbers using the well-known linear congruential algorithm and 48-bit integer arithmetic.</p> <p>Functions <code>drand48()</code> and <code>erand48()</code> return non-negative double-precision floating-point values uniformly distributed over the interval [0.0, 1.0).</p> <p>Functions <code>lrnd48()</code> and <code>nrand48()</code> return non-negative long integers uniformly distributed over the interval $[0, 2^{31}]$.</p> <p>Functions <code>mrnd48()</code> and <code>jrnd48()</code> return signed long integers uniformly distributed over the interval $[-2^{31}, 2^{31}]$.</p> <p>Functions <code>srand48()</code>, <code>seed48()</code>, and <code>lcong48()</code> are initialization entry points, one of which should be invoked before either <code>drand48()</code>, <code>lrnd48()</code>, or <code>mrnd48()</code> is called. (Although it is not recommended practice, constant default initializer values will be supplied automatically if <code>drand48()</code>, <code>lrnd48()</code>, or <code>mrnd48()</code> is called without a prior call to an initialization entry point.) Functions <code>erand48()</code>, <code>nrand48()</code>, and <code>jrnd48()</code> do not require an initialization entry point to be called first.</p> <p>All the routines work by generating a sequence of 48-bit integer values, X_i, according to the linear congruential formula</p> $X_{n+1} = (aX_n + c) \bmod m \quad n \geq 0.$ <p>The parameter $m = 2^{48}$; hence 48-bit integer arithmetic is performed. Unless <code>lcong48()</code> has been invoked, the multiplier value a and the addend value c are given by</p> $a = 5DEECE66D_{16} = 273673163155_8$

$c = B_{16} = 13_8$.

The value returned by any of the functions `drand48()`, `erand48()`, `lrand48()`, `nrand48()`, `mrnd48()`, or `jrand48()` is computed by first generating the next 48-bit X_i in the sequence. Then the appropriate number of bits, according to the type of data item to be returned, are copied from the high-order (leftmost) bits of X_i and transformed into the returned value.

The functions `drand48()`, `lrand48()`, and `mrnd48()` store the last 48-bit X_i generated in an internal buffer. X_i must be initialized prior to being invoked. The functions `erand48()`, `nrand48()`, and `jrand48()` require the calling program to provide storage for the successive X_i values in the array specified as an argument when the functions are invoked. These routines do not have to be initialized; the calling program must place the desired initial value of X_i into the array and pass it as an argument. By using different arguments, functions `erand48()`, `nrand48()`, and `jrand48()` allow separate modules of a large program to generate several *independent* streams of pseudo-random numbers, that is, the sequence of numbers in each stream will *not* depend upon how many times the routines have been called to generate numbers for the other streams.

The initializer function `srand48()` sets the high-order 32 bits of X_i to the 32 bits contained in its argument. The low-order 16 bits of X_i are set to the arbitrary value $330E_{16}$.

The initializer function `seed48()` sets the value of X_i to the 48-bit value specified in the argument array. In addition, the previous value of X_i is copied into a 48-bit internal buffer, used only by `seed48()`, and a pointer to this buffer is the value returned by `seed48()`. This returned pointer, which can just be ignored if not needed, is useful if a program is to be restarted from a given point at some future time — use the pointer to get at and store the last X_i value, and then use this value to reinitialize using `seed48()` when the program is restarted.

The initialization function `lcng48()` allows the user to specify the initial X_i the multiplier value a , and the addend value c . Argument array elements `param[0-2]` specify X_i , `param[3-5]` specify the multiplier a , and `param[6]` specifies the 16-bit addend c . After `lcng48()` has been called, a subsequent call to either `srand48()` or `seed48()` will restore the “standard” multiplier and addend values, a and c , specified above.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Safe

SEE ALSO `rand(3C)`, `attributes(5)`, `standards(5)`

dup2(3C)

NAME	dup2 – duplicate an open file descriptor								
SYNOPSIS	<pre>#include <unistd.h> int dup2 (int <i>fdes</i>, int <i>fdes2</i>);</pre>								
DESCRIPTION	<p>The dup2 () function causes the file descriptor <i>fdes2</i> to refer to the same file as <i>fdes</i>. The <i>fdes</i> argument is a file descriptor referring to an open file, and <i>fdes2</i> is a non-negative integer less than the current value for the maximum number of open file descriptors allowed the calling process. See getrlimit(2). If <i>fdes2</i> already refers to an open file, not <i>fdes</i>, it is closed first. If <i>fdes2</i> refers to <i>fdes</i>, or if <i>fdes</i> is not a valid open file descriptor, <i>fdes2</i> will not be closed first.</p> <p>The dup2 () function is equivalent to fcntl (<i>fdes</i>, F_DUP2FD, <i>fdes2</i>).</p>								
RETURN VALUES	Upon successful completion a non-negative integer representing the file descriptor is returned. Otherwise, -1 is returned and errno is set to indicate the error.								
ERRORS	<p>The dup2 () function will fail if:</p> <table><tr><td>EBADF</td><td>The <i>fdes</i> argument is not a valid open file descriptor.</td></tr><tr><td>EBADF</td><td>The <i>fdes2</i> argument is negative or is not less than the current resource limit returned by getrlimit (RLIMIT_NOFILE, . . .).</td></tr><tr><td>EINTR</td><td>A signal was caught during the dup2 () call.</td></tr><tr><td>EMFILE</td><td>The process has too many open files. See fcntl(2).</td></tr></table>	EBADF	The <i>fdes</i> argument is not a valid open file descriptor.	EBADF	The <i>fdes2</i> argument is negative or is not less than the current resource limit returned by getrlimit (RLIMIT_NOFILE, . . .).	EINTR	A signal was caught during the dup2 () call.	EMFILE	The process has too many open files. See fcntl(2).
EBADF	The <i>fdes</i> argument is not a valid open file descriptor.								
EBADF	The <i>fdes2</i> argument is negative or is not less than the current resource limit returned by getrlimit (RLIMIT_NOFILE, . . .).								
EINTR	A signal was caught during the dup2 () call.								
EMFILE	The process has too many open files. See fcntl(2).								
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>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>Async-Signal-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Async-Signal-Safe		
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Interface Stability	Standard								
MT-Level	Async-Signal-Safe								
SEE ALSO	close(2), creat(2), exec(2), fcntl(2), getrlimit(2), open(2), pipe(2), lockf(3C), attributes(5), standards(5)								

NAME	econvert, fconvert, gconvert, seconvert, sfconvert, sgconvert, qeconvert, qfconvert, qgconvert – output conversion
SYNOPSIS	<pre>#include <floatingpoint.h> char *econvert(double value, int ndigit, int *decpt, int *sign, char *buf) ; char *fconvert(double value, int ndigit, int *decpt, int *sign, char *buf) ; char *gconvert(double value, int ndigit, int trailing, char *buf) ; char *seconvert(single *value, int ndigit, int *decpt, int *sign, char *buf) ; char *sfconvert(single *value, int ndigit, int *decpt, int *sign, char *buf) ; char *sgconvert(single *value, int ndigit, int trailing, char *buf) ; char *qeconvert(quadruple *value, int ndigit, int *decpt, int *sign, char *buf) ; char *qfconvert(quadruple *value, int ndigit, int *decpt, int *sign, char *buf) ; char *qgconvert(quadruple *value, int ndigit, int trailing, char *buf) ;</pre>
DESCRIPTION	<p>The <code>econvert()</code> function converts the <i>value</i> to a null-terminated string of <i>ndigit</i> ASCII digits in <i>buf</i> and returns a pointer to <i>buf</i>. <i>buf</i> should contain at least <i>ndigit</i>+1 characters. The position of the decimal point relative to the beginning of the string is stored indirectly through <i>decpt</i>. Thus <i>buf</i> == "314" and <i>*decpt</i> == 1 corresponds to the numerical value 3.14, while <i>buf</i> == "314" and <i>*decpt</i> == -1 corresponds to the numerical value .0314. If the sign of the result is negative, the word pointed to by <i>sign</i> is nonzero; otherwise it is zero. The least significant digit is rounded.</p> <p>The <code>fconvert()</code> function works much like <code>econvert()</code>, except that the correct digit has been rounded as if for <code>printf(%w.nf)</code> output with $n=ndigit$ digits to the right of the decimal point. <i>ndigit</i> can be negative to indicate rounding to the left of the decimal point. The return value is a pointer to <i>buf</i>. <i>buf</i> should contain at least $310+max(0,ndigit)$ characters to accommodate any double-precision <i>value</i>.</p> <p>The <code>gconvert()</code> function converts the <i>value</i> to a null-terminated ASCII string in <i>buf</i> and returns a pointer to <i>buf</i>. It produces <i>ndigit</i> significant digits in fixed-decimal format, like <code>printf(%w.nf)</code>, if possible, and otherwise in floating-decimal format, like <code>printf(%w.ne)</code>; in either case <i>buf</i> is ready for printing, with sign and exponent. The result corresponds to that obtained by</p> <pre>(void) printf(buf, ``%w.ng'`, value) ;</pre> <p>If <i>trailing</i> = 0, trailing zeros and a trailing point are suppressed, as in <code>printf(%g)</code>. If <i>trailing</i> != 0, trailing zeros and a trailing point are retained, as in <code>printf(%#g)</code>.</p>

econvert(3C)

The `seconvert()`, `sfconvert()`, and `sgconvert()` functions are single-precision versions of these functions, and are more efficient than the corresponding double-precision versions. A pointer rather than the value itself is passed to avoid C's usual conversion of single-precision arguments to double.

The `geconvert()`, `qfconvert()`, and `ggconvert()` functions are quadruple-precision versions of these functions. The `qfconvert()` function can overflow the `decimal_record` field `ds` if `value` is too large. In that case, `buf[0]` is set to zero.

The `ecvt()`, `fcvt()` and `gcvt()` functions are versions of `econvert()`, `fconvert()`, and `gconvert()`, respectively, that are documented on the [ecvt\(3C\)](#) manual page. They constitute the default implementation of these functions and conform to the X/Open CAE Specification, System Interfaces and Headers, Issue 4, Version 2.

USAGE IEEE Infinities and NaNs are treated similarly by these functions. "NaN" is returned for NaN, and "Inf" or "Infinity" for Infinity. The longer form is produced when `ndigit` \geq 8.

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [ecvt\(3C\)](#), [sprintf\(3C\)](#), [attributes\(5\)](#)

NAME	ecvt, fcvt, gcvt – convert floating-point number to string
SYNOPSIS	<pre>#include <stdlib.h> char *ecvt(double value, int ndigit, int *restrict decpt, int *restrict sign); char *fcvt(double value, int ndigit, int *restrict decpt, int *restrict sign); char *gcvt(double value, int ndigit, char *buf);</pre>
DESCRIPTION	<p>The <code>ecvt()</code>, <code>fcvt()</code> and <code>gcvt()</code> functions convert floating-point numbers to null-terminated strings.</p> <p>ecvt() The <code>ecvt()</code> function converts <i>value</i> to a null-terminated string of <i>ndigit</i> digits (where <i>ndigit</i> is reduced to an unspecified limit determined by the precision of a <code>double</code>) and returns a pointer to the string. The high-order digit is non-zero, unless the value is 0. The low-order digit is rounded. The position of the radix character relative to the beginning of the string is stored in the integer pointed to by <i>decpt</i> (negative means to the left of the returned digits). The radix character is not included in the returned string. If the sign of the result is negative, the integer pointed to by <i>sign</i> is non-zero, otherwise it is 0.</p> <p>If the converted value is out of range or is not representable, the contents of the returned string are unspecified.</p> <p>fcvt() The <code>fcvt()</code> function is identical to <code>ecvt()</code> except that <i>ndigit</i> specifies the number of digits desired after the radix point. The total number of digits in the result string is restricted to an unspecified limit as determined by the precision of a <code>double</code>.</p> <p>gcvt() The <code>gcvt()</code> function converts <i>value</i> to a null-terminated string (similar to that of the <code>%g</code> format of <code>printf(3C)</code>) in the array pointed to by <i>buf</i> and returns <i>buf</i>. It produces <i>ndigit</i> significant digits (limited to an unspecified value determined by the precision of a <code>double</code>) in <code>%f</code> if possible, or <code>%e</code> (scientific notation) otherwise. A minus sign is included in the returned string if <i>value</i> is less than 0. A radix character is included in the returned string if <i>value</i> is not a whole number. Trailing zeros are suppressed where <i>value</i> is not a whole number. The radix character is determined by the current locale. If <code>setlocale(3C)</code> has not been called successfully, the default locale, <code>POSIX</code>, is used. The default locale specifies a period (<code>.</code>) as the radix character. The <code>LC_NUMERIC</code> category determines the value of the radix character within the current locale.</p>
RETURN VALUES	<p>The <code>ecvt()</code> and <code>fcvt()</code> functions return a pointer to a null-terminated string of digits.</p> <p>The <code>gcvt()</code> function returns <i>buf</i>.</p>
ERRORS	No errors are defined.
USAGE	The return values from <code>ecvt()</code> and <code>fcvt()</code> might point to thread-specific data that can be overwritten by subsequent calls to these functions by the same thread.

ecvt(3C)

For portability to implementations conforming to earlier versions of Solaris, [sprintf\(3C\)](#) is preferred over this function.

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Safe

SEE ALSO [printf\(3C\)](#), [setlocale\(3C\)](#), [sprintf\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	encrypt – encoding function						
SYNOPSIS	<pre>#include <crypt.h> void encrypt(char <i>block</i>[64], int <i>edflag</i>);</pre>						
Standard conforming	<pre>#include <unistd.h> void encrypt(char <i>block</i>[64], int <i>edflag</i>);</pre>						
DESCRIPTION	<p>The <code>encrypt()</code> function provides (rather primitive) access to the hashing algorithm employed by the <code>crypt(3C)</code> function. The key generated by <code>setkey(3C)</code> is used to encrypt the string <code>block</code> with <code>encrypt()</code>.</p> <p>The <code>block</code> argument to <code>encrypt()</code> is an array of length 64 bytes containing only the bytes with numerical value of 0 and 1. The array is modified in place to a similar array using the key set by <code>setkey(3C)</code>. If <code>edflag</code> is 0, the argument is encoded. If <code>edflag</code> is 1, the argument may be decoded (see the USAGE section below); if the argument is not decoded, <code>errno</code> will be set to <code>ENOSYS</code>.</p>						
RETURN VALUES	The <code>encrypt()</code> function returns no value.						
ERRORS	<p>The <code>encrypt()</code> function will fail if:</p> <p><code>ENOSYS</code> The functionality is not supported on this implementation.</p>						
USAGE	<p>In some environments, decoding may not be implemented. This is related to U.S. Government restrictions on encryption and decryption routines: the DES decryption algorithm cannot be exported outside the U.S.A. Historical practice has been to ship a different version of the encryption library without the decryption feature in the routines supplied. Thus the exported version of <code>encrypt()</code> does encoding but not decoding.</p> <p>Because <code>encrypt()</code> does not return a value, applications wishing to check for errors should set <code>errno</code> to 0, call <code>encrypt()</code>, then test <code>errno</code> and, if it is non-zero, assume an error has occurred.</p>						
ATTRIBUTES	<p>See <code>attributes(5)</code> 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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	<code>crypt(3C)</code> , <code>setkey(3C)</code> , <code>attributes(5)</code>						

end(3C)

NAME	end, _end, etext, _etext, edata, _edata – last locations in program
SYNOPSIS	<pre>extern int _etext; extern int _edata; extern int _end;</pre>
DESCRIPTION	<p>These names refer neither to routines nor to locations with interesting contents; only their addresses are meaningful.</p> <p><code>_etext</code> The address of <code>_etext</code> is the first location after the program text.</p> <p><code>_edata</code> The address of <code>_edata</code> is the first location after the initialized data region.</p> <p><code>_end</code> The address of <code>_end</code> is the first location after the uninitialized data region.</p>
USAGE	<p>When execution begins, the program break (the first location beyond the data) coincides with <code>_end</code>, but the program break may be reset by the <code>brk(2)</code>, <code>malloc(3C)</code>, and the standard input/output library (see <code>stdio(3C)</code>), functions by the profile (<code>-p</code>) option of <code>cc(1B)</code>, and so on. Thus, the current value of the program break should be determined by <code>sbrk ((char *) 0)</code>.</p> <p>References to <code>end</code>, <code>etext</code>, and <code>edata</code>, without a preceding underscore will be aliased to the associated symbol that begins with the underscore.</p>
SEE ALSO	<code>cc(1B)</code> , <code>brk(2)</code> , <code>malloc(3C)</code> , <code>stdio(3C)</code>

NAME	euclen, euccol, eucscol – get byte length and display width of EUC characters				
SYNOPSIS	<pre>#include <euc.h> int euclen(const unsigned char *s); int euccol(const unsigned char *s); int eucscol(const unsigned char *str);</pre>				
DESCRIPTION	<p>The <code>euclen()</code> function returns the length in bytes of the Extended Unix Code (EUC) character pointed to by <code>s</code>, including single-shift characters, if present.</p> <p>The <code>euccol()</code> function returns the screen column width of the EUC character pointed to by <code>s</code>.</p> <p>The <code>eucscol()</code> function returns the screen column width of the EUC string pointed to by <code>str</code>.</p> <p>For the <code>euclen()</code> and <code>euccol()</code>, functions, <code>s</code> points to the first byte of the character. This byte is examined to determine its codeset. The character type table for the current <i>locale</i> is used for codeset byte length and display width information.</p>				
USAGE	<p>These functions will work only with EUC locales.</p> <p>These functions can be used safely in multithreaded applications, as long as <code>setlocale(3C)</code> is not called to change the locale.</p>				
ATTRIBUTES	See <code>attributes(5)</code> 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>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe with exceptions				
SEE ALSO	<code>getwidth(3C)</code> , <code>setlocale(3C)</code> , <code>attributes(5)</code>				

exit(3C)

NAME exit, _exithandle – terminate process

SYNOPSIS

```
#include <stdlib.h>

void exit(int status);

void _exithandle(void);
```

DESCRIPTION The `exit()` function terminates a process by calling first `_exithandle()` and then `_exit()` (see `exit(2)`).

The `_exithandle()` function calls any functions registered through the `atexit(3C)` function in the reverse order of their registration. This action includes executing all finalization code from the `.fini` sections of all objects that are part of the process.

The `_exithandle()` function is intended for use *only* with `_exit()`, and allows for specialized processing such as `dldump(3C)` to be performed. Normal process execution should not be continued after a call to `_exithandle()` has occurred, as internal data structures may have been torn down due to `atexit()` or `.fini` processing.

The symbols `EXIT_SUCCESS` and `EXIT_FAILURE` are defined in the header `<stdlib.h>` and may be used as the value of `status` to indicate successful or unsuccessful termination, respectively.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `exit(2)`, `atexit(3C)`, `dldump(3C)`, `attributes(5)`, `standards(5)`

NAME	fattach – attach a STREAMS- or doors-based file descriptor to an object in the file system name space																		
SYNOPSIS	<pre>#include <stropts.h> int fattach(int <i>fildev</i>, const char *<i>path</i>);</pre>																		
DESCRIPTION	<p>The <code>fattach()</code> function attaches a STREAMS- or doors-based file descriptor to an object in the file system name space, effectively associating a name with <i>fildev</i>. The <i>fildev</i> argument must be a valid open file descriptor representing a STREAMS or doors file. The <i>path</i> argument is a path name of an existing object and the user must have appropriate privileges or be the owner of the file and have write permissions. All subsequent operations on <i>path</i> will operate on the STREAMS or doors file until the STREAMS or doors file is detached from the node. The <i>fildev</i> argument can be attached to more than one <i>path</i>, that is, a stream or door can have several names associated with it.</p> <p>The attributes of the named stream or door (see <code>stat(2)</code>), are initialized as follows: the permissions, user ID, group ID, and times are set to those of <i>path</i>, the number of links is set to 1, and the size and device identifier are set to those of the streams or doors device associated with <i>fildev</i>. If any attributes of the named stream or door are subsequently changed (for example, <code>chmod(2)</code>), the attributes of the underlying object are not affected.</p>																		
RETURN VALUES	Upon successful completion, <code>fattach()</code> returns 0. Otherwise it returns -1 and sets <code>errno</code> to indicate an error.																		
ERRORS	<p>The <code>fattach()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EACCES</td> <td>The user is the owner of <i>path</i> but does not have write permissions on <i>path</i> or <i>fildev</i> is locked.</td> </tr> <tr> <td>EBADF</td> <td>The <i>fildev</i> argument is not a valid open file descriptor.</td> </tr> <tr> <td>EBUSY</td> <td>The <i>path</i> argument is currently a mount point or has a STREAMS or doors file descriptor attached to it.</td> </tr> <tr> <td>EINVAL</td> <td>The <i>path</i> argument is a file in a remotely mounted directory.</td> </tr> <tr> <td>EINVAL</td> <td>The <i>fildev</i> argument does not represent a STREAMS or doors file.</td> </tr> <tr> <td>ELOOP</td> <td>Too many symbolic links were encountered in translating <i>path</i>.</td> </tr> <tr> <td>ENAMETOOLONG</td> <td>The size of <i>path</i> exceeds <code>{PATH_MAX}</code>, or the component of a path name is longer than <code>{NAME_MAX}</code> while <code>{_POSIX_NO_TRUNC}</code> is in effect.</td> </tr> <tr> <td>ENOENT</td> <td>The <i>path</i> argument does not exist.</td> </tr> <tr> <td>ENOTDIR</td> <td>A component of a path prefix is not a directory.</td> </tr> </table>	EACCES	The user is the owner of <i>path</i> but does not have write permissions on <i>path</i> or <i>fildev</i> is locked.	EBADF	The <i>fildev</i> argument is not a valid open file descriptor.	EBUSY	The <i>path</i> argument is currently a mount point or has a STREAMS or doors file descriptor attached to it.	EINVAL	The <i>path</i> argument is a file in a remotely mounted directory.	EINVAL	The <i>fildev</i> argument does not represent a STREAMS or doors file.	ELOOP	Too many symbolic links were encountered in translating <i>path</i> .	ENAMETOOLONG	The size of <i>path</i> exceeds <code>{PATH_MAX}</code> , or the component of a path name is longer than <code>{NAME_MAX}</code> while <code>{_POSIX_NO_TRUNC}</code> is in effect.	ENOENT	The <i>path</i> argument does not exist.	ENOTDIR	A component of a path prefix is not a directory.
EACCES	The user is the owner of <i>path</i> but does not have write permissions on <i>path</i> or <i>fildev</i> is locked.																		
EBADF	The <i>fildev</i> argument is not a valid open file descriptor.																		
EBUSY	The <i>path</i> argument is currently a mount point or has a STREAMS or doors file descriptor attached to it.																		
EINVAL	The <i>path</i> argument is a file in a remotely mounted directory.																		
EINVAL	The <i>fildev</i> argument does not represent a STREAMS or doors file.																		
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .																		
ENAMETOOLONG	The size of <i>path</i> exceeds <code>{PATH_MAX}</code> , or the component of a path name is longer than <code>{NAME_MAX}</code> while <code>{_POSIX_NO_TRUNC}</code> is in effect.																		
ENOENT	The <i>path</i> argument does not exist.																		
ENOTDIR	A component of a path prefix is not a directory.																		

fattach(3C)

EPERM

The effective user ID is not the owner of *path* or a user with the appropriate privileges.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO

`fdetach(1M)`, `chmod(2)`, `mount(2)`, `stat(2)`, `door_create(3DOOR)`, `fdetach(3C)`, `isastream(3C)`, `attributes(5)`, `standards(5)`, `streamio(7I)`

STREAMS Programming Guide

NAME	__fbufsize, __flbf, __fpending, __fpurge, __freadable, __freading, __fsetlocking, __fwritable, __fwriting, _flushlbf – interfaces to stdio FILE structure
SYNOPSIS	<pre>#include <stdio.h> #include <stdio_ext.h> size_t __fbufsiz(FILE *stream); int __flbf(FILE *stream); size_t __fpending(FILE *stream); void __fpurge(FILE *stream); int __freadable(FILE *stream); int __freading(FILE *stream); int __fsetlocking(FILE *stream, int type); int __fwritable(FILE *stream); int __fwriting(FILE *stream); void _flushlbf(void);</pre>
DESCRIPTION	<p>These functions provide portable access to the members of the stdio(3C) FILE structure.</p> <p>The <code>__fbufsize()</code> function returns in bytes the size of the buffer currently in use by the given stream.</p> <p>The <code>__flbf()</code> function returns non-zero if the stream is line-buffered.</p> <p>The <code>__fpending</code> function returns in bytes the amount of output pending on a stream.</p> <p>The <code>__fpurge()</code> function discards any pending buffered I/O on the stream.</p> <p>The <code>__freadable()</code> function returns non-zero if it is possible to read from a stream.</p> <p>The <code>__freading()</code> function returns non-zero if the file is open readonly, or if the last operation on the stream was a read operation such as fread(3C) or fgetc(3C). Otherwise it returns 0.</p> <p>The <code>__fsetlocking()</code> function allows the type of locking performed by <code>stdio</code> on a given stream to be controlled by the programmer.</p> <p>If <code>type</code> is <code>FSETLOCKING_INTERNAL</code>, <code>stdio</code> performs implicit locking around every operation on the given stream. This is the default system behavior on that stream.</p>

__fbufsize(3C)

If *type* is FSETLOCKING_BYCALLER, `stdio` assumes that the caller is responsible for maintaining the integrity of the stream in the face of access by multiple threads. If there is only one thread accessing the stream, nothing further needs to be done. If multiple threads are accessing the stream, then the caller can use the `flockfile()`, `funlockfile()`, and `ftrylockfile()` functions described on the [flockfile\(3C\)](#) manual page to provide the appropriate locking. In both this and the case where *type* is FSETLOCKING_INTERNAL, `__fsetlocking()` returns the previous state of the stream.

If *type* is FSETLOCKING_QUERY, `__fsetlocking()` returns the current state of the stream without changing it.

The `__fwritable()` function returns non-zero if it is possible to write on a stream.

The `__fwriting()` function returns non-zero if the file is open write-only or append-only, or if the last operation on the stream was a write operation such as [fwrite\(3C\)](#) or [fputc\(3C\)](#). Otherwise it returns 0.

The `_flushlbf()` function flushes all line-buffered files. It is used when reading from a line-buffered file.

USAGE

Although the contents of the `stdio` FILE structure have always been private to the `stdio` implementation, some applications have needed to obtain information about a `stdio` stream that was not accessible through a supported interface. These applications have resorted to accessing fields of the FILE structure directly, rendering them possibly non-portable to new implementations of `stdio`, or more likely, preventing enhancements to `stdio` that would cause those applications to break.

In the 64-bit environment, the FILE structure is opaque. The functions described here are provided as a means of obtaining the information that up to now has been retrieved directly from the FILE structure. Because they are based on the needs of existing applications (such as `mh` and `emacs`), they may be extended as other programs are ported. Although they may still be non-portable to other operating systems, they will be compatible from each Solaris release to the next. Interfaces that are more portable are under development.

ATTRIBUTES

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	<code>__fsetlocking()</code> is Unsafe; all others are MT-Safe
Interface Stability	Evolving

SEE ALSO

[fgetc\(3C\)](#), [flockfile\(3C\)](#), [fputc\(3C\)](#), [fread\(3C\)](#), [fwrite\(3C\)](#), [stdio\(3C\)](#), [attributes\(5\)](#)

NAME	fclose – close a stream																
SYNOPSIS	<pre>#include <stdio.h> int fclose(FILE *<i>stream</i>);</pre>																
DESCRIPTION	<p>The <code>fclose()</code> function causes the stream pointed to by <i>stream</i> to be flushed and the associated file to be closed. Any unwritten buffered data for the stream is written to the file; any unread buffered data is discarded. The stream is disassociated from the file. If the associated buffer was automatically allocated, it is deallocated.</p> <p>The <code>fclose()</code> function marks for update the <code>st_ctime</code> and <code>st_mtime</code> fields of the underlying file if the stream is writable and if buffered data has not yet been written to the file. It will perform a <code>close(2)</code> operation on the file descriptor that is associated with the stream pointed to by <i>stream</i>.</p> <p>After the call to <code>fclose()</code>, any use of <i>stream</i> causes undefined behavior.</p> <p>The <code>fclose()</code> function is performed automatically for all open files upon calling <code>exit(2)</code>.</p>																
RETURN VALUES	Upon successful completion, <code>fclose()</code> returns 0. Otherwise, it returns EOF and sets <code>errno</code> to indicate the error.																
ERRORS	<p>The <code>fclose()</code> function will fail if:</p> <table border="0"> <tr> <td style="vertical-align: top;">EAGAIN</td> <td>The <code>O_NONBLOCK</code> flag is set for the file descriptor underlying <i>stream</i> and the process would be delayed in the write operation.</td> </tr> <tr> <td style="vertical-align: top;">EBADF</td> <td>The file descriptor underlying stream is not valid.</td> </tr> <tr> <td style="vertical-align: top;">EFBIG</td> <td>An attempt was made to write a file that exceeds the maximum file size or the process's file size limit; or the file is a regular file and an attempt was made to write at or beyond the offset maximum associated with the corresponding stream.</td> </tr> <tr> <td style="vertical-align: top;">EINTR</td> <td>The <code>fclose()</code> function was interrupted by a signal.</td> </tr> <tr> <td style="vertical-align: top;">EIO</td> <td>The process is a member of a background process group attempting to write to its controlling terminal, <code>TOSTOP</code> is set, the process is neither ignoring nor blocking <code>SIGTTOU</code> and the process group of the process is orphaned.</td> </tr> <tr> <td style="vertical-align: top;">ENOSPC</td> <td>There was no free space remaining on the device containing the file.</td> </tr> <tr> <td style="vertical-align: top;">EPIPE</td> <td>An attempt is made to write to a pipe or FIFO that is not open for reading by any process. A <code>SIGPIPE</code> signal will also be sent to the calling thread.</td> </tr> </table> <p>The <code>fclose()</code> function may fail if:</p> <table border="0"> <tr> <td style="vertical-align: top;">ENXIO</td> <td>A request was made of a non-existent device, or the request was beyond the limits of the device.</td> </tr> </table>	EAGAIN	The <code>O_NONBLOCK</code> flag is set for the file descriptor underlying <i>stream</i> and the process would be delayed in the write operation.	EBADF	The file descriptor underlying stream is not valid.	EFBIG	An attempt was made to write a file that exceeds the maximum file size or the process's file size limit; or the file is a regular file and an attempt was made to write at or beyond the offset maximum associated with the corresponding stream.	EINTR	The <code>fclose()</code> function was interrupted by a signal.	EIO	The process is a member of a background process group attempting to write to its controlling terminal, <code>TOSTOP</code> is set, the process is neither ignoring nor blocking <code>SIGTTOU</code> and the process group of the process is orphaned.	ENOSPC	There was no free space remaining on the device containing the file.	EPIPE	An attempt is made to write to a pipe or FIFO that is not open for reading by any process. A <code>SIGPIPE</code> signal will also be sent to the calling thread.	ENXIO	A request was made of a non-existent device, or the request was beyond the limits of the device.
EAGAIN	The <code>O_NONBLOCK</code> flag is set for the file descriptor underlying <i>stream</i> and the process would be delayed in the write operation.																
EBADF	The file descriptor underlying stream is not valid.																
EFBIG	An attempt was made to write a file that exceeds the maximum file size or the process's file size limit; or the file is a regular file and an attempt was made to write at or beyond the offset maximum associated with the corresponding stream.																
EINTR	The <code>fclose()</code> function was interrupted by a signal.																
EIO	The process is a member of a background process group attempting to write to its controlling terminal, <code>TOSTOP</code> is set, the process is neither ignoring nor blocking <code>SIGTTOU</code> and the process group of the process is orphaned.																
ENOSPC	There was no free space remaining on the device containing the file.																
EPIPE	An attempt is made to write to a pipe or FIFO that is not open for reading by any process. A <code>SIGPIPE</code> signal will also be sent to the calling thread.																
ENXIO	A request was made of a non-existent device, or the request was beyond the limits of the device.																

fclose(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `close(2)`, `exit(2)`, `getrlimit(2)`, `ulimit(2)`, `fopen(3C)`, `stdio(3C)`, `attributes(5)`, `standards(5)`

NAME	fdetach – detach a name from a STREAMS-based file descriptor																
SYNOPSIS	<pre>#include <stropts.h> int fdetach(const char *path);</pre>																
DESCRIPTION	<p>The <code>fdetach()</code> function detaches a STREAMS-based file from the file to which it was attached by a previous call to <code>fattach(3C)</code>. The <i>path</i> argument points to the pathname of the attached STREAMS file. The process must have appropriate privileges or be the owner of the file. A successful call to <code>fdetach()</code> causes all pathnames that named the attached STREAMS file to again name the file to which the STREAMS file was attached. All subsequent operations on <i>path</i> will operate on the underlying file and not on the STREAMS file.</p> <p>All open file descriptions established while the STREAMS file was attached to the file referenced by <i>path</i>, will still refer to the STREAMS file after the <code>fdetach()</code> has taken effect.</p> <p>If there are no open file descriptors or other references to the STREAMS file, then a successful call to <code>fdetach()</code> has the same effect as performing the last <code>close(2)</code> on the attached file.</p>																
RETURN VALUES	Upon successful completion, <code>fdetach()</code> returns 0. Otherwise, it returns -1 and sets <code>errno</code> to indicate the error.																
ERRORS	<p>The <code>fdetach()</code> function will fail if:</p> <table border="0"> <tr> <td style="vertical-align: top;">EACCES</td> <td>Search permission is denied on a component of the path prefix.</td> </tr> <tr> <td style="vertical-align: top;">EPERM</td> <td>The effective user ID is not the owner of <i>path</i> and the process does not have appropriate privileges.</td> </tr> <tr> <td style="vertical-align: top;">ENOTDIR</td> <td>A component of the path prefix is not a directory.</td> </tr> <tr> <td style="vertical-align: top;">ENOENT</td> <td>A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.</td> </tr> <tr> <td style="vertical-align: top;">EINVAL</td> <td>The <i>path</i> argument names a file that is not currently attached.</td> </tr> <tr> <td style="vertical-align: top;">ENAMETOOLONG</td> <td>The size of a pathname 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.</td> </tr> <tr> <td style="vertical-align: top;">ELOOP</td> <td>Too many symbolic links were encountered in resolving <i>path</i>.</td> </tr> </table> <p>The <code>fdetach()</code> function may fail if:</p> <table border="0"> <tr> <td style="vertical-align: top;">ENAMETOOLONG</td> <td>Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>PATH_MAX</code>.</td> </tr> </table>	EACCES	Search permission is denied on a component of the path prefix.	EPERM	The effective user ID is not the owner of <i>path</i> and the process does not have appropriate privileges.	ENOTDIR	A component of the path prefix is not a directory.	ENOENT	A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.	EINVAL	The <i>path</i> argument names a file that is not currently attached.	ENAMETOOLONG	The size of a pathname 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.	ELOOP	Too many symbolic links were encountered in resolving <i>path</i> .	ENAMETOOLONG	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>PATH_MAX</code> .
EACCES	Search permission is denied on a component of the path prefix.																
EPERM	The effective user ID is not the owner of <i>path</i> and the process does not have appropriate privileges.																
ENOTDIR	A component of the path prefix is not a directory.																
ENOENT	A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.																
EINVAL	The <i>path</i> argument names a file that is not currently attached.																
ENAMETOOLONG	The size of a pathname 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.																
ELOOP	Too many symbolic links were encountered in resolving <i>path</i> .																
ENAMETOOLONG	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>PATH_MAX</code> .																

fdetach(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

SEE ALSO `fdetach(1M)`, `close(2)`, `fattach(3C)`, `attributes(5)`, `standards(5)`, `streamio(7I)`

STREAMS Programming Guide

NAME	fdopen – associate a stream with a file descriptor												
SYNOPSIS	<pre>#include <stdio.h> FILE *fdopen(int <i>fildes</i>, const char *<i>mode</i>);</pre>												
DESCRIPTION	<p>The <code>fdopen()</code> function associates a stream with a file descriptor <i>fildes</i>.</p> <p>The <i>mode</i> argument is a character string having one of the following values:</p> <table border="0"> <tr> <td style="padding-right: 20px;">r or rb</td> <td>Open a file for reading.</td> </tr> <tr> <td>w or wb</td> <td>Open a file for writing.</td> </tr> <tr> <td>a or ab</td> <td>Open a file for writing at end of file.</td> </tr> <tr> <td>r+ or rb+ or r+b</td> <td>Open a file for update (reading and writing).</td> </tr> <tr> <td>w+ or wb+ or w+b</td> <td>Open a file for update (reading and writing).</td> </tr> <tr> <td>a+ or ab+ or a+b</td> <td>Open a file for update (reading and writing) at end of file.</td> </tr> </table> <p>The meaning of these flags is exactly as specified for the fopen(3C) function, except that modes beginning with <i>w</i> do not cause truncation of the file.</p> <p>The mode of the stream must be allowed by the file access mode of the open file. The file position indicator associated with the new stream is set to the position indicated by the file offset associated with the file descriptor.</p> <p>The <code>fdopen()</code> function preserves the offset maximum previously set for the open file description corresponding to <i>fildes</i>.</p> <p>The error and end-of-file indicators for the stream are cleared. The <code>fdopen()</code> function may cause the <code>st_atime</code> field of the underlying file to be marked for update.</p> <p>If <i>fildes</i> refers to a shared memory object, the result of the <code>fdopen()</code> function is unspecified.</p>	r or rb	Open a file for reading.	w or wb	Open a file for writing.	a or ab	Open a file for writing at end of file.	r+ or rb+ or r+b	Open a file for update (reading and writing).	w+ or wb+ or w+b	Open a file for update (reading and writing).	a+ or ab+ or a+b	Open a file for update (reading and writing) at end of file.
r or rb	Open a file for reading.												
w or wb	Open a file for writing.												
a or ab	Open a file for writing at end of file.												
r+ or rb+ or r+b	Open a file for update (reading and writing).												
w+ or wb+ or w+b	Open a file for update (reading and writing).												
a+ or ab+ or a+b	Open a file for update (reading and writing) at end of file.												
RETURN VALUES	<p>Upon successful completion, <code>fdopen()</code> returns a pointer to a stream. Otherwise, a null pointer is returned and <code>errno</code> is set to indicate the error.</p> <p>The <code>fdopen()</code> function may fail and not set <code>errno</code> if there are no free <code>stdio</code> streams.</p>												
ERRORS	<p>The <code>fdopen()</code> function may fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EBADF</td> <td>The <i>fildes</i> argument is not a valid file descriptor.</td> </tr> <tr> <td>EINVAL</td> <td>The <i>mode</i> argument is not a valid mode.</td> </tr> <tr> <td>EMFILE</td> <td>The number of streams currently open in the calling process is either <code>FOPEN_MAX</code> or <code>STREAM_MAX</code>.</td> </tr> </table>	EBADF	The <i>fildes</i> argument is not a valid file descriptor.	EINVAL	The <i>mode</i> argument is not a valid mode.	EMFILE	The number of streams currently open in the calling process is either <code>FOPEN_MAX</code> or <code>STREAM_MAX</code> .						
EBADF	The <i>fildes</i> argument is not a valid file descriptor.												
EINVAL	The <i>mode</i> argument is not a valid mode.												
EMFILE	The number of streams currently open in the calling process is either <code>FOPEN_MAX</code> or <code>STREAM_MAX</code> .												

fdopen(3C)

ENOMEM Insufficient space to allocate a buffer.

USAGE The minimum number of streams that a process can have open at one time is `FOPEN_MAX`. For 32-bit applications, however, the underlying ABIs require that no file descriptor used to access the file underlying an `stdio` `STREAM` have a value greater than 255.

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

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

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

NAME | ferror, feof, clearerr, fileno – stream status inquiries

SYNOPSIS | `#include <stdio.h>`
`int ferror(FILE *stream);`
`int feof(FILE *stream);`
`void clearerr(FILE *stream);`
`int fileno(FILE *stream);`

DESCRIPTION | The `ferror()` function returns a non-zero value when an error has previously occurred reading from or writing to the named *stream* (see [intro\(3\)](#)). It returns 0 otherwise.

The `feof()` function returns a non-zero value when EOF has previously been detected reading the named input *stream*. It returns 0 otherwise.

The `clearerr()` function resets the error indicator and EOF indicator to 0 on the named *stream*.

The `fileno()` function returns the integer file descriptor associated with the named *stream*; see [open\(2\)](#).

ATTRIBUTES | See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO | [open\(2\)](#), [intro\(3\)](#), [fopen\(3C\)](#), [stdio\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

fflush(3C)

NAME	fflush – flush a stream														
SYNOPSIS	<pre>#include <stdio.h> int fflush(FILE *stream);</pre>														
DESCRIPTION	<p>If <i>stream</i> points to an output stream or an update stream in which the most recent operation was not input, <code>fflush()</code> causes any unwritten data for that stream to be written to the file, and the <code>st_ctime</code> and <code>st_mtime</code> fields of the underlying file are marked for update.</p> <p>If <i>stream</i> points to an input stream or an update stream into which the most recent operation was input, that stream is flushed if it is seekable and is not already at end-of-file. Flushing an input stream discards any buffered input and adjusts the file pointer such that the next input operation accesses the byte after the last one read. A stream is seekable if the underlying file is not a pipe, FIFO, socket, or TTY device.</p> <p>If <i>stream</i> is a null pointer, <code>fflush()</code> performs this flushing action on all streams for which the behavior is defined above.</p> <p>An input stream, seekable or non-seekable, can be flushed by explicitly calling <code>fflush()</code> with a non-null argument specifying that stream.</p>														
RETURN VALUES	Upon successful completion, <code>fflush()</code> returns 0. Otherwise, it returns EOF and sets <code>errno</code> to indicate the error.														
ERRORS	<p>The <code>fflush()</code> function will fail if:</p> <table><tr><td>EAGAIN</td><td>The <code>O_NONBLOCK</code> flag is set for the file descriptor underlying <i>stream</i> and the process would be delayed in the write operation.</td></tr><tr><td>EBADF</td><td>The file descriptor underlying <i>stream</i> is not valid.</td></tr><tr><td>EFBIG</td><td>An attempt was made to write a file that exceeds the maximum file size or the process's file size limit; or the file is a regular file and an attempt was made to write at or beyond the offset maximum associated with the corresponding stream.</td></tr><tr><td>EINTR</td><td>The <code>fflush()</code> function was interrupted by a signal.</td></tr><tr><td>EIO</td><td>The process is a member of a background process group attempting to write to its controlling terminal, <code>TOSTOP</code> is set, the process is neither ignoring nor blocking <code>SIGTTOU</code>, and the process group of the process is orphaned.</td></tr><tr><td>ENOSPC</td><td>There was no free space remaining on the device containing the file.</td></tr><tr><td>EPIPE</td><td>An attempt is made to write to a pipe or FIFO that is not open for reading by any process. A <code>SIGPIPE</code> signal will also be sent to the calling process.</td></tr></table> <p>The <code>fflush()</code> function may fail if:</p>	EAGAIN	The <code>O_NONBLOCK</code> flag is set for the file descriptor underlying <i>stream</i> and the process would be delayed in the write operation.	EBADF	The file descriptor underlying <i>stream</i> is not valid.	EFBIG	An attempt was made to write a file that exceeds the maximum file size or the process's file size limit; or the file is a regular file and an attempt was made to write at or beyond the offset maximum associated with the corresponding stream.	EINTR	The <code>fflush()</code> function was interrupted by a signal.	EIO	The process is a member of a background process group attempting to write to its controlling terminal, <code>TOSTOP</code> is set, the process is neither ignoring nor blocking <code>SIGTTOU</code> , and the process group of the process is orphaned.	ENOSPC	There was no free space remaining on the device containing the file.	EPIPE	An attempt is made to write to a pipe or FIFO that is not open for reading by any process. A <code>SIGPIPE</code> signal will also be sent to the calling process.
EAGAIN	The <code>O_NONBLOCK</code> flag is set for the file descriptor underlying <i>stream</i> and the process would be delayed in the write operation.														
EBADF	The file descriptor underlying <i>stream</i> is not valid.														
EFBIG	An attempt was made to write a file that exceeds the maximum file size or the process's file size limit; or the file is a regular file and an attempt was made to write at or beyond the offset maximum associated with the corresponding stream.														
EINTR	The <code>fflush()</code> function was interrupted by a signal.														
EIO	The process is a member of a background process group attempting to write to its controlling terminal, <code>TOSTOP</code> is set, the process is neither ignoring nor blocking <code>SIGTTOU</code> , and the process group of the process is orphaned.														
ENOSPC	There was no free space remaining on the device containing the file.														
EPIPE	An attempt is made to write to a pipe or FIFO that is not open for reading by any process. A <code>SIGPIPE</code> signal will also be sent to the calling process.														

fflush(3C)

ENXIO A request was made of a non-existent device, or the request was beyond the limits of the device.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO getrlimit(2), ulimit(2), attributes(5), standards(5)

ffs(3C)

NAME ffs – find first set bit

SYNOPSIS `#include <strings.h>`
`int ffs(const int i);`

DESCRIPTION The `ffs()` function finds the first bit set (beginning with the least significant bit) and returns the index of that bit. Bits are numbered starting at one (the least significant bit).

RETURN VALUES The `ffs()` function returns the index of the first bit set. If *i* is 0, then `ffs()` returns 0.

ERRORS No errors are defined.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `attributes(5)`, `standards(5)`

NAME	fgetc, getc, getc_unlocked, getchar, getchar_unlocked, getw – get a byte from a stream
SYNOPSIS	<pre>#include <stdio.h> int fgetc(FILE *stream); int getc(FILE *stream); int getc_unlocked(FILE *stream); int getchar(void); int getchar_unlocked(void); int getw(FILE *stream);</pre>
DESCRIPTION	<p>The <code>fgetc()</code> function obtains the next byte (if present) as an unsigned char converted to an <code>int</code>, from the input stream pointed to by <code>stream</code>, and advances the associated file position indicator for the stream (if defined).</p> <p>For standard-conforming (see standards(5)) applications, if the end-of-file indicator for the stream is set, <code>fgetc()</code> returns EOF whether or not a next byte is present.</p> <p>The <code>fgetc()</code> function may mark the <code>st_atime</code> field of the file associated with <code>stream</code> for update. The <code>st_atime</code> field will be marked for update by the first successful execution of <code>fgetc()</code>, fgets(3C), fread(3C), fscanf(3C), <code>getc()</code>, <code>getchar()</code>, gets(3C) or scanf(3C) using <code>stream</code> that returns data not supplied by a prior call to ungetc(3C) or ungetwc(3C).</p> <p>The <code>getc()</code> function is functionally identical to <code>fgetc()</code>, except that it is implemented as a macro. It runs faster than <code>fgetc()</code>, but it takes up more space per invocation and its name cannot be passed as an argument to a function call.</p> <p>The <code>getchar()</code> routine is equivalent to <code>getc(stdin)</code>. It is implemented as a macro.</p> <p>The <code>getc_unlocked()</code> and <code>getchar_unlocked()</code> routines are variants of <code>getc()</code> and <code>getchar()</code>, respectively, that do not lock the stream. It is the caller's responsibility to acquire the stream lock before calling these routines and releasing the lock afterwards; see flockfile(3C) and stdio(3C). These routines are implemented as macros.</p> <p>The <code>getw()</code> function reads the next word from the <code>stream</code>. The size of a word is the size of an <code>int</code> and may vary from environment to environment. The <code>getw()</code> function presumes no special alignment in the file.</p> <p>The <code>getw()</code> function may mark the <code>st_atime</code> field of the file associated with <code>stream</code> for update. The <code>st_atime</code> field will be marked for update by the first successful execution of <code>fgetc()</code>, fgets(3C), fread(3C), <code>getc()</code>, <code>getchar()</code>, gets(3C), fscanf(3C) or scanf(3C) using <code>stream</code> that returns data not supplied by a prior call to ungetc(3C).</p>
RETURN VALUES	Upon successful completion, <code>fgetc()</code> , <code>getc()</code> , <code>getc_unlocked()</code> , <code>getchar()</code> , <code>getchar_unlocked()</code> , and <code>getw()</code> return the next byte from the input stream pointed to by <code>stream</code> . If the stream is at end-of-file, the end-of-file indicator for the

fgetc(3C)

stream is set and these functions return EOF. For standard-conforming (see standards(5)) applications, if the end-of-file indicator for the stream is set, these functions return EOF whether or not the stream is at end-of-file. If a read error occurs, the error indicator for the stream is set, EOF is returned, and `errno` is set to indicate the error.

ERRORS The `fgetc()`, `getc()`, `getc_unlocked()`, `getchar()`, `getchar_unlocked()`, and `getw()` functions will fail if data needs to be read and:

EAGAIN	The <code>O_NONBLOCK</code> flag is set for the file descriptor underlying <i>stream</i> and the process would be delayed in the <code>fgetc()</code> operation.
EBADF	The file descriptor underlying <i>stream</i> is not a valid file descriptor open for reading.
EINTR	The read operation was terminated due to the receipt of a signal, and no data was transferred.
EIO	A physical I/O error has occurred, or the process is in a background process group attempting to read from its controlling terminal, and either the process is ignoring or blocking the <code>SIGTTIN</code> signal or the process group is orphaned. This error may also be generated for implementation-dependent reasons.
EOVERFLOW	The file is a regular file and an attempt was made to read at or beyond the offset maximum associated with the corresponding stream.

The `fgetc()`, `getc()`, `getc_unlocked()`, `getchar()`, `getchar_unlocked()`, and `getw()` functions may fail if:

ENOMEM	Insufficient storage space is available.
ENXIO	A request was made of a non-existent device, or the request was outside the capabilities of the device.

USAGE If the integer value returned by `fgetc()`, `getc()`, `getc_unlocked()`, `getchar()`, `getchar_unlocked()`, and `getw()` is stored into a variable of type `char` and then compared against the integer constant `EOF`, the comparison may never succeed, because sign-extension of a variable of type `char` on widening to integer is implementation-dependent.

The `ferror(3C)` or `feof(3C)` functions must be used to distinguish between an error condition and an end-of-file condition.

Functions exist for the `getc()`, `getc_unlocked()`, `getchar()`, and `getchar_unlocked()` macros. To get the function form, the macro name must be undefined (for example, `#undef getc`).

When the macro forms are used, `getc()` and `getc_unlocked()` evaluate the *stream* argument more than once. In particular, `getc(*f++)`; does not work sensibly. The `fgetc()` function should be used instead when evaluating the *stream* argument has side effects.

Because of possible differences in word length and byte ordering, files written using `getw()` are machine-dependent, and may not be read using `getw()` on a different processor.

The `getw()` function is inherently byte stream-oriented and is not tenable in the context of either multibyte character streams or wide-character streams. Application programmers are recommended to use one of the character-based input functions instead.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	<code>fgetc()</code> , <code>getc()</code> , <code>getc_unlocked()</code> , <code>getchar()</code> , and <code>getchar_unlocked()</code> are Standard.
MT-Level	See <code>NOTES</code> below.

SEE ALSO `intro(3)`, `__fsetlocking(3C)`, `fclose(3C)`, `feof(3C)`, `fgets(3C)`, `fgetwc(3C)`, `fgetws(3C)`, `flockfile(3C)`, `fopen(3C)`, `fread(3C)`, `fscanf(3C)`, `gets(3C)`, `putc(3C)`, `scanf(3C)`, `stdio(3C)`, `ungetc(3C)`, `ungetwc(3C)`, `attributes(5)`, `standards(5)`

NOTES The `fgetc()`, `getc()`, `getchar()`, and `getw()` routines are MT-Safe in multithreaded applications. The `getc_unlocked()` and `getchar_unlocked()` routines are unsafe in multithreaded applications.

fgetpos(3C)

NAME	fgetpos – get current file position information				
SYNOPSIS	<pre>#include <stdio.h> int fgetpos(FILE *stream, fpos_t *pos);</pre>				
DESCRIPTION	The <code>fgetpos()</code> function stores the current value of the file position indicator for the stream pointed to by <i>stream</i> in the object pointed to by <i>pos</i> . The value stored contains unspecified information usable by <code>fsetpos(3C)</code> for repositioning the stream to its position at the time of the call to <code>fgetpos()</code> .				
RETURN VALUES	Upon successful completion, <code>fgetpos()</code> returns 0. Otherwise, it returns a non-zero value and sets <code>errno</code> to indicate the error.				
ERRORS	The <code>fgetpos()</code> function may fail if: EBADF The file descriptor underlying <i>stream</i> is not valid. ESPIPE The file descriptor underlying <i>stream</i> is associated with a pipe, a FIFO, or a socket. EOVERFLOW The current value of the file position cannot be represented correctly in an object of type <code>fpos_t</code> .				
USAGE	The <code>fgetpos()</code> function has a transitional interface for 64-bit file offsets. See <code>lf64(5)</code> .				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Standard				
SEE ALSO	<code>fopen(3C)</code> , <code>fsetpos(3C)</code> , <code>ftell(3C)</code> , <code>rewind(3C)</code> , <code>ungetc(3C)</code> , <code>attributes(5)</code> , <code>lf64(5)</code> , <code>standards(5)</code>				

NAME	fgetwc – get a wide-character code from a stream										
SYNOPSIS	<pre>#include <stdio.h> #include <wchar.h> wint_t fgetwc (FILE*stream) ;</pre>										
DESCRIPTION	<p>The <code>fgetwc()</code> function obtains the next character (if present) from the input stream pointed to by <i>stream</i>, converts that to the corresponding wide-character code and advances the associated file position indicator for the stream (if defined).</p> <p>If an error occurs, the resulting value of the file position indicator for the stream is indeterminate.</p> <p>The <code>fgetwc()</code> function may mark the <code>st_atime</code> field of the file associated with <i>stream</i> for update. The <code>st_atime</code> field will be marked for update by the first successful execution of <code>fgetwc()</code>, <code>fgetc(3C)</code>, <code>fgets(3C)</code>, <code>fgetws(3C)</code>, <code>fread(3C)</code>, <code>fscanf(3C)</code>, <code>getc(3C)</code>, <code>getchar(3C)</code>, <code>gets(3C)</code>, or <code>scanf(3C)</code> using <i>stream</i> that returns data not supplied by a prior call to <code>ungetc(3C)</code> or <code>ungetwc(3C)</code>.</p>										
RETURN VALUES	<p>Upon successful completion the <code>fgetwc()</code> function returns the wide-character code of the character read from the input stream pointed to by <i>stream</i> converted to a type <code>wint_t</code>.</p> <p>For standard-conforming (see <code>standards(5)</code>) applications, if the end-of-file indicator for the stream is set, <code>fgetwc()</code> returns <code>WEOF</code> whether or not the stream is at end-of-file.</p> <p>If a read error occurs, the error indicator for the stream is set, <code>fgetwc()</code> returns <code>WEOF</code> and sets <code>errno</code> to indicate the error.</p> <p>If an encoding error occurs, the error indicator for the stream is set, <code>fgetwc()</code> returns <code>WEOF</code>, and <code>errno</code> is set to indicate the error.</p>										
ERRORS	<p>The <code>fgetwc()</code> function will fail if data needs to be read and:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>EAGAIN</code></td> <td>The <code>O_NONBLOCK</code> flag is set for the file descriptor underlying <i>stream</i> and the process would be delayed in the <code>fgetwc()</code> operation.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>EBADF</code></td> <td>The file descriptor underlying <i>stream</i> is not a valid file descriptor open for reading.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>EINTR</code></td> <td>The read operation was terminated due to the receipt of a signal, and no data was transferred.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>EIO</code></td> <td>A physical I/O error has occurred, or the process is in a background process group attempting to read from its controlling terminal and either the process is ignoring or blocking the <code>SIGTTIN</code> signal or the process group is orphaned.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>EOVERFLOW</code></td> <td>The file is a regular file and an attempt was made to read at or beyond the offset maximum associated with the corresponding</td> </tr> </table>	<code>EAGAIN</code>	The <code>O_NONBLOCK</code> flag is set for the file descriptor underlying <i>stream</i> and the process would be delayed in the <code>fgetwc()</code> operation.	<code>EBADF</code>	The file descriptor underlying <i>stream</i> is not a valid file descriptor open for reading.	<code>EINTR</code>	The read operation was terminated due to the receipt of a signal, and no data was transferred.	<code>EIO</code>	A physical I/O error has occurred, or the process is in a background process group attempting to read from its controlling terminal and either the process is ignoring or blocking the <code>SIGTTIN</code> signal or the process group is orphaned.	<code>EOVERFLOW</code>	The file is a regular file and an attempt was made to read at or beyond the offset maximum associated with the corresponding
<code>EAGAIN</code>	The <code>O_NONBLOCK</code> flag is set for the file descriptor underlying <i>stream</i> and the process would be delayed in the <code>fgetwc()</code> operation.										
<code>EBADF</code>	The file descriptor underlying <i>stream</i> is not a valid file descriptor open for reading.										
<code>EINTR</code>	The read operation was terminated due to the receipt of a signal, and no data was transferred.										
<code>EIO</code>	A physical I/O error has occurred, or the process is in a background process group attempting to read from its controlling terminal and either the process is ignoring or blocking the <code>SIGTTIN</code> signal or the process group is orphaned.										
<code>EOVERFLOW</code>	The file is a regular file and an attempt was made to read at or beyond the offset maximum associated with the corresponding										

fgetwc(3C)

stream.

The `fgetwc()` function may fail if:

ENOMEM	Insufficient storage space is available.
ENXIO	A request was made of a non-existent device, or the request was outside the capabilities of the device.
EILSEQ	The data obtained from the input stream does not form a valid character.

USAGE The `ferror(3C)` or `feof(3C)` functions must be used to distinguish between an error condition and an end-of-file condition.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

SEE ALSO `feof(3C)`, `ferror(3C)`, `fgetc(3C)`, `fgets(3C)`, `fgetws(3C)`, `fopen(3C)`, `fread(3C)`, `fscanf(3C)`, `getc(3C)`, `getchar(3C)`, `gets(3C)`, `scanf(3C)`, `setlocale(3C)`, `ungetc(3C)`, `ungetwc(3C)`, `attributes(5)`, `standards(5)`

NAME	floating_to_decimal, single_to_decimal, double_to_decimal, extended_to_decimal, quadruple_to_decimal – convert floating-point value to decimal record
SYNOPSIS	<pre>#include <floatingpoint.h> void single_to_decimal(single *px, decimal_mode *pm, decimal_record *pd, fp_exception_field_type *ps); void double_to_decimal(double *px, decimal_mode *pm, decimal_record *pd, fp_exception_field_type *ps); void extended_to_decimal(extended *px, decimal_mode *pm, decimal_record *pd, fp_exception_field_type *ps); void quadruple_to_decimal(quadruple *px, decimal_mode *pm, decimal_record *pd, fp_exception_field_type *ps);</pre>
DESCRIPTION	<p>The <code>floating_to_decimal()</code> functions convert the floating-point value at <code>*px</code> into a decimal record at <code>*pd</code>, observing the modes specified in <code>*pm</code> and setting exceptions in <code>*ps</code>. If there are no IEEE exceptions, <code>*ps</code> will be zero.</p> <p>If <code>*px</code> is zero, infinity, or NaN, then only <code>pd->sign</code> and <code>pd->fpclass</code> are set. Otherwise <code>pd->exponent</code> and <code>pd->ds</code> are also set so that</p> <p>$(\text{sig}) * (\text{pd->ds}) * 10^{(\text{pd->exponent})}$ is a correctly rounded approximation to <code>*px</code>, where <code>sig</code> is +1 or -1, depending upon whether <code>pd->sign</code> is 0 or -1. <code>pd->ds</code> has at least one and no more than <code>DECIMAL_STRING_LENGTH-1</code> significant digits because one character is used to terminate the string with a NULL.</p> <p><code>pd->ds</code> is correctly rounded according to the IEEE rounding modes in <code>pm->rd</code>. <code>*ps</code> has <code>fp_inexact</code> set if the result was inexact, and has <code>fp_overflow</code> set if the string result does not fit in <code>pd->ds</code> because of the limitation <code>DECIMAL_STRING_LENGTH</code>.</p> <p>If <code>pm->df == floating_form</code>, then <code>pd->ds</code> always contains <code>pm->ndigits</code> significant digits. Thus if <code>*px == 12.34</code> and <code>pm->ndigits == 8</code>, then <code>pd->ds</code> will contain 12340000 and <code>pd->exponent</code> will contain -6.</p> <p>If <code>pm->df == fixed_form</code> and <code>pm->ndigits >= 0</code>, then <code>pd->ds</code> always contains <code>pm->ndigits</code> after the point and as many digits as necessary before the point. Since the latter is not known in advance, the total number of digits required is returned in <code>pd->ndigits</code>; if that number <code>>= DECIMAL_STRING_LENGTH</code>, then <code>ds</code> is undefined. <code>pd->exponent</code> always gets <code>-pm->ndigits</code>. Thus if <code>*px == 12.34</code> and <code>pm->ndigits == 1</code>, then <code>pd->ds</code> gets 123, <code>pd->exponent</code> gets -1, and <code>pd->ndigits</code> gets 3.</p> <p>If <code>pm->df == fixed_form</code> and <code>pm->ndigits < 0</code>, then <code>pd->ds</code> always contains <code>-pm->ndigits</code> trailing zeros; in other words, rounding occurs <code>-pm->ndigits</code> to the left of the decimal point, but the digits rounded away are retained as zeros. The total number of digits required is in <code>pd->ndigits</code>. <code>pd->exponent</code> always gets 0. Thus if <code>*px == 12.34</code> and <code>pm->ndigits == -1</code>, then <code>pd->ds</code> gets 10, <code>pd->exponent</code> gets 0, and <code>pd->ndigits</code> gets 2.</p> <p><code>pd->more</code> is not used.</p>

floating_to_decimal(3C)

[econvert\(3C\)](#), [fconvert\(3C\)](#), [gconvert\(3C\)](#), [printf\(3C\)](#), and [sprintf\(3C\)](#) all use `double_to_decimal()`.

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [econvert\(3C\)](#), [fconvert\(3C\)](#), [gconvert\(3C\)](#), [printf\(3C\)](#), [sprintf\(3C\)](#), [attributes\(5\)](#)

NAME	flock – apply or remove an advisory lock on an open file
SYNOPSIS	<pre> /usr/ucb/cc[flag ...] file ... #include <sys/file.h> int flock(fd, operation); int fd, operation; </pre>
DESCRIPTION	<p>flock() applies or removes an <i>advisory</i> lock on the file associated with the file descriptor <i>fd</i>. The compatibility version of flock() has been implemented on top of fcntl(2) locking. It does not provide complete binary compatibility.</p> <p>Advisory locks allow cooperating processes to perform consistent operations on files, but do not guarantee exclusive access (that is, processes may still access files without using advisory locks, possibly resulting in inconsistencies).</p> <p>The locking mechanism allows two types of locks: shared locks and exclusive locks. More than one process may hold a shared lock for a file at any given time, but multiple exclusive, or both shared and exclusive, locks may not exist simultaneously on a file.</p> <p>A lock is applied by specifying an <i>operation</i> parameter LOCK_SH for a shared lock or LOCK_EX for an exclusive lock. The <i>operation</i> parameter may be ORed with LOCK_NB to make the operation non-blocking. To unlock an existing lock, the <i>operation</i> should be LOCK_UN.</p> <p>Read permission is required on a file to obtain a shared lock, and write permission is required to obtain an exclusive lock. Locking a segment that is already locked by the calling process causes the old lock type to be removed and the new lock type to take effect.</p> <p>Requesting a lock on an object that is already locked normally causes the caller to block until the lock may be acquired. If LOCK_NB is included in <i>operation</i>, then this will not happen; instead, the call will fail and the error EWOULDBLOCK will be returned.</p>
RETURN VALUES	<p>flock() returns:</p> <p>0 on success.</p> <p>-1 on failure and sets errno to indicate the error.</p>
ERRORS	<p>EBADF The argument <i>fd</i> is an invalid descriptor.</p> <p>EINVAL <i>operation</i> is not a valid argument.</p> <p>EOPNOTSUPP The argument <i>fd</i> refers to an object other than a file.</p> <p>EWOULDBLOCK The file is locked and the LOCK_NB option was specified.</p>
SEE ALSO	lockd(1M), chmod(2), close(2), dup(2), exec(2), fcntl(2), fork(2), open(2), lockf(3C)

flock(3UCB)

NOTES Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.

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

Processes blocked awaiting a lock may be awakened by signals.

Mandatory locking may occur, depending on the mode bits of the file. See `chmod(2)`.

Locks obtained through the `flock()` mechanism under SunOS 4.1 were known only within the system on which they were placed. This is no longer true.

NAME	flockfile, funlockfile, ftrylockfile – acquire and release stream lock
SYNOPSIS	<pre>#include <stdio.h> void flockfile(FILE *stream) ; void funlockfile(FILE *stream) ; int ftrylockfile(FILE *stream) ;</pre>
DESCRIPTION	<p>The flockfile() function acquires an internal lock of a stream <i>stream</i>. If the lock is already acquired by another thread, the thread calling flockfile() is suspended until it can acquire the lock. In the case that the stream lock is available, flockfile() not only acquires the lock, but keeps track of the number of times it is being called by the current thread. This implies that the stream lock can be acquired more than once by the same thread.</p> <p>The funlockfile() function releases the lock being held by the current thread. In the case of recursive locking, this function must be called the same number of times flockfile() was called. After the number of funlockfile() calls is equal to the number of flockfile() calls, the stream lock is available for other threads to acquire.</p> <p>The ftrylockfile() function acquires an internal lock of a stream <i>stream</i>, only if that object is available. In essence ftrylockfile() is a non-blocking version of flockfile().</p>
RETURN VALUES	The ftrylockfile() function returns 0 on success and non-zero to indicate a lock cannot be acquired.
EXAMPLES	<p>EXAMPLE 1 A sample program of flockfile().</p> <p>The following example prints everything out together, blocking other threads that might want to write to the same file between calls to fprintf(3C):</p> <pre>FILE iop; flockfile(iop); fprintf(iop, "hello "); fprintf(iop, "world"); fputc(iop, 'a'); funlockfile(iop);</pre> <p>An unlocked interface is available in case performance is an issue. For example:</p> <pre>flockfile(iop); while (!feof(iop)) { *c++ = getc_unlocked(iop); } funlockfile(iop);</pre>

flockfile(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `intro(3)`, `__fsetlocking(3C)`, `ferror(3C)`, `fprintf(3C)`, `getc(3C)`, `putc(3C)`, `stdio(3C)`, `ungetc(3C)`, `attributes(5)`, `standards(5)`

NOTES The interfaces on this page are as specified in IEEE Std 1003.1:2001. See `standards(5)`.

NAME	fmtmsg – display a message on stderr or system console
SYNOPSIS	<pre>#include <fmtmsg.h> int fmtmsg(long <i>classification</i>, const char *<i>label</i>, int <i>severity</i>, const char *<i>text</i>, const char *<i>action</i>, const char *<i>tag</i>);</pre>
DESCRIPTION	<p>The <code>fmtmsg()</code> function writes a formatted message to <code>stderr</code>, to the console, or to both, on a message's classification component. It can be used instead of the traditional <code>printf(3C)</code> interface to display messages to <code>stderr</code>, and in conjunction with <code>gettext(3C)</code>, provides a simple interface for producing language-independent applications.</p> <p>A formatted message consists of up to five standard components (<i>label</i>, <i>severity</i>, <i>text</i>, <i>action</i>, and <i>tag</i>) as described below. The <i>classification</i> component is not part of the standard message displayed to the user, but rather defines the source of the message and directs the display of the formatted message.</p> <p><i>classification</i> Contains identifiers from the following groups of major classifications and subclassifications. Any one identifier from a subclass may be used in combination by ORing the values together with a single identifier from a different subclass. Two or more identifiers from the same subclass should not be used together, with the exception of identifiers from the display subclass. (Both display subclass identifiers may be used so that messages can be displayed to both <code>stderr</code> and the system console).</p> <ul style="list-style-type: none"> ■ “Major classifications” identify the source of the condition. Identifiers are: <code>MM_HARD</code> (hardware), <code>MM_SOFT</code> (software), and <code>MM_FIRM</code> (firmware). ■ “Message source subclassifications” identify the type of software in which the problem is spotted. Identifiers are: <code>MM_APPL</code> (application), <code>MM_UTIL</code> (utility), and <code>MM_OP SYS</code> (operating system). ■ “Display subclassifications” indicate where the message is to be displayed. Identifiers are: <code>MM_PRINT</code> to display the message on the standard error stream, <code>MM_CONSOLE</code> to display the message on the system console. Neither, either, or both identifiers may be used. ■ “Status subclassifications” indicate whether the application will recover from the condition. Identifiers are: <code>MM_RECOVER</code> (recoverable) and <code>MM_NRECOV</code> (non-recoverable). ■ An additional identifier, <code>MM_NULLMC</code>, indicates that no classification component is supplied for the message. <p><i>label</i> Identifies the source of the message. The format of this component is two fields separated by a colon. The first field is up to 10 characters long; the second is up to 14 characters. Suggested usage</p>

fmtmsg(3C)

Environment Variables	<i>severity</i>	<p>is that <i>label</i> identifies the package in which the application resides as well as the program or application name. For example, the <i>label</i> <code>UX:cat</code> indicates the UNIX System V package and the <code>cat(1)</code> utility.</p> <p>Indicates the seriousness of the condition. Identifiers for the standard levels of <i>severity</i> are:</p> <ul style="list-style-type: none"> ■ <code>MM_HALT</code> indicates that the application has encountered a severe fault and is halting. Produces the print string <code>HALT</code>. ■ <code>MM_ERROR</code> indicates that the application has detected a fault. Produces the print string <code>ERROR</code>. ■ <code>MM_WARNING</code> indicates a condition out of the ordinary that might be a problem and should be watched. Produces the print string <code>WARNING</code>. ■ <code>MM_INFO</code> provides information about a condition that is not in error. Produces the print string <code>INFO</code>. ■ <code>MM_NOSEV</code> indicates that no severity level is supplied for the message. <p>Other severity levels may be added by using the <code>addseverity()</code> routine.</p>
	<i>text</i>	<p>Describes the condition that produced the message. The <i>text</i> string is not limited to a specific size.</p>
	<i>action</i>	<p>Describes the first step to be taken in the error recovery process. <code>fmtmsg()</code> precedes each action string with the prefix: <code>TOFIX: .</code> The <i>action</i> string is not limited to a specific size.</p>
	<i>tag</i>	<p>An identifier which references on-line documentation for the message. Suggested usage is that <i>tag</i> includes the <i>label</i> and a unique identifying number. A sample <i>tag</i> is <code>UX:cat:146</code>.</p>
	<p>The <code>MSGVERB</code> and <code>SEV_LEVEL</code> environment variables control the behavior of <code>fmtmsg()</code> as follows:</p>	
<code>MSGVERB</code>	<p>This variable determines which message components <code>fmtmsg()</code> selects when writing messages to <code>stderr</code>. Its value is a colon-separated list of optional keywords and can be set as follows:</p> <pre>MSGVERB=[keyword[:keyword[: . . .]]] export MSGVERB</pre> <p>Valid <i>keywords</i> are: <i>label</i>, <i>severity</i>, <i>text</i>, <i>action</i>, and <i>tag</i>. If <code>MSGVERB</code> contains a keyword for a component and the component's value is not the component's null value, <code>fmtmsg()</code> includes that component in the message when writing the message to <code>stderr</code>. If <code>MSGVERB</code> does not include a keyword for a message component, that component is not included in the display of the</p>	

message. The keywords may appear in any order. If MSGVERB is not defined, if its value is the null string, if its value is not of the correct format, or if it contains keywords other than the valid ones listed above, `fmtmsg()` selects all components.

The first time `fmtmsg()` is called, it examines MSGVERB to determine which message components are to be selected when generating a message to write to the standard error stream, `stderr`. The values accepted on the initial call are saved for future calls.

The MSGVERB environment variable affects only those components that are selected for display to the standard error stream. All message components are included in console messages.

SEV_LEVEL

This variable defines severity levels and associates print strings with them for use by `fmtmsg()`. The standard severity levels listed below cannot be modified. Additional severity levels can also be defined, redefined, and removed using `addseverity()` (see [addseverity\(3C\)](#)). If the same severity level is defined by both SEV_LEVEL and `addseverity()`, the definition by `addseverity()` takes precedence.

0	(no severity is used)
1	HALT
2	ERROR
3	WARNING
4	INFO

The SEV_LEVEL variable can be set as follows:

```
SEV_LEVEL=[description[:description[: . . .]]]
export SEV_LEVEL
```

where *description* is a comma-separated list containing three fields:

description=severity_keyword,level,printstring

The *severity_keyword* field is a character string that is used as the keyword on the `-s severity` option to the `fmtmsg(1)` utility. (This field is not used by the `fmtmsg()` function.)

The *level* field is a character string that evaluates to a positive integer (other than 0, 1, 2, 3, or 4, which are reserved for the standard severity levels). If the keyword *severity_keyword* is used, *level* is the severity value passed on to the `fmtmsg()` function.

fmtmsg(3C)

The *printstring* field is the character string used by `fmtmsg()` in the standard message format whenever the severity value *level* is used.

If a *description* in the colon list is not a three-field comma list, or if the second field of a comma list does not evaluate to a positive integer, that *description* in the colon list is ignored.

The first time `fmtmsg()` is called, it examines the `SEV_LEVEL` environment variable, if defined, to determine whether the environment expands the levels of severity beyond the five standard levels and those defined using `addseverity()`. The values accepted on the initial call are saved for future calls.

Use in Applications

One or more message components may be systematically omitted from messages generated by an application by using the null value of the argument for that component.

The table below indicates the null values and identifiers for `fmtmsg()` arguments.

Argument	Type	Null-Value	Identifier
<i>label</i>	char*	(char*) NULL	MM_NULLLBL
<i>severity</i>	int	0	MM_NULLSEV
<i>class</i>	long	0L	MM_NULLMC
<i>text</i>	char*	(char*) NULL	MM_NULLTXT
<i>action</i>	char*	(char*) NULL	MM_NULLACT
<i>tag</i>	char*	(char*) NULL	MM_NULLTAG

Another means of systematically omitting a component is by omitting the component keyword(s) when defining the `MSGVERB` environment variable (see the `Environment Variables` section above).

RETURN VALUES

The `fmtmsg()` returns the following values:

MM_OK	The function succeeded.
MM_NOTOK	The function failed completely.
MM_NOMSG	The function was unable to generate a message on the standard error stream, but otherwise succeeded.
MM_NOCON	The function was unable to generate a console message, but otherwise succeeded.

EXAMPLES

EXAMPLE 1 The following example of `fmtmsg()`:

```
fmtmsg(MM_PRINT, "UX:cat", MM_ERROR, "invalid syntax",
"refer to manual", "UX:cat:001")
```

produces a complete message in the standard message format:

```
UX:cat: ERROR: invalid syntax
TO FIX: refer to manual    UX:cat:001
```

EXAMPLE 2 When the environment variable `MSGVERB` is set as follows:

```
MSGVERB=severity:text:action
```

and the Example 1 is used, `fmtmsg()` produces:

```
ERROR: invalid syntax
TO FIX: refer to manual
```

EXAMPLE 3 When the environment variable `SEV_LEVEL` is set as follows:

```
SEV_LEVEL=note,5,NOTE
```

the following call to `fmtmsg()`

```
fmtmsg(MM_UTIL | MM_PRINT, "UX:cat", 5, "invalid syntax",
"refer to manual", "UX:cat:001")
```

produces

```
UX:cat: NOTE: invalid syntax
TO FIX: refer to manual    UX:cat:001
```

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Safe

SEE ALSO

`fmtmsg(1)`, `addseverity(3C)`, `gettext(3C)`, `printf(3C)`, `attributes(5)`, `standards(5)`

fnmatch(3C)

NAME	fnmatch – match filename or path name														
SYNOPSIS	<pre>#include <fnmatch.h> int fnmatch(const char *<i>pattern</i>, const char *<i>string</i>, int <i>flags</i>);</pre>														
DESCRIPTION	<p>The <code>fnmatch()</code> function matches patterns as described on the <code>fnmatch(5)</code> manual page. It checks the <i>string</i> argument to see if it matches the <i>pattern</i> argument.</p> <p>The <i>flags</i> argument modifies the interpretation of <i>pattern</i> and <i>string</i>. It is the bitwise inclusive OR of zero or more of the following flags defined in the header <code><fnmatch.h></code>.</p> <table><tr><td><code>FNM_PATHNAME</code></td><td>If set, a slash (/) character in <i>string</i> will be explicitly matched by a slash in <i>pattern</i>; it will not be matched by either the asterisk (*) or question-mark (?) special characters, nor by a bracket ([]) expression.</td></tr><tr><td></td><td>If not set, the slash character is treated as an ordinary character.</td></tr><tr><td><code>FNM_NOESCAPE</code></td><td>If not set, a backslash character (\) in <i>pattern</i> followed by any other character will match that second character in <i>string</i>. In particular, "\\\" will match a backslash in <i>string</i>.</td></tr><tr><td></td><td>If set, a backslash character will be treated as an ordinary character.</td></tr><tr><td><code>FNM_PERIOD</code></td><td>If set, a leading period in <i>string</i> will match a period in <i>pattern</i>; where the location of "leading" is indicated by the value of <code>FNM_PATHNAME</code>:</td></tr><tr><td></td><td><ul style="list-style-type: none">■ If <code>FNM_PATHNAME</code> is set, a period is "leading" if it is the first character in <i>string</i> or if it immediately follows a slash.■ If <code>FNM_PATHNAME</code> is not set, a period is "leading" only if it is the first character of <i>string</i>.</td></tr><tr><td></td><td>If not set, no special restrictions are placed on matching a period.</td></tr></table>	<code>FNM_PATHNAME</code>	If set, a slash (/) character in <i>string</i> will be explicitly matched by a slash in <i>pattern</i> ; it will not be matched by either the asterisk (*) or question-mark (?) special characters, nor by a bracket ([]) expression.		If not set, the slash character is treated as an ordinary character.	<code>FNM_NOESCAPE</code>	If not set, a backslash character (\) in <i>pattern</i> followed by any other character will match that second character in <i>string</i> . In particular, "\\\" will match a backslash in <i>string</i> .		If set, a backslash character will be treated as an ordinary character.	<code>FNM_PERIOD</code>	If set, a leading period in <i>string</i> will match a period in <i>pattern</i> ; where the location of "leading" is indicated by the value of <code>FNM_PATHNAME</code> :		<ul style="list-style-type: none">■ If <code>FNM_PATHNAME</code> is set, a period is "leading" if it is the first character in <i>string</i> or if it immediately follows a slash.■ If <code>FNM_PATHNAME</code> is not set, a period is "leading" only if it is the first character of <i>string</i>.		If not set, no special restrictions are placed on matching a period.
<code>FNM_PATHNAME</code>	If set, a slash (/) character in <i>string</i> will be explicitly matched by a slash in <i>pattern</i> ; it will not be matched by either the asterisk (*) or question-mark (?) special characters, nor by a bracket ([]) expression.														
	If not set, the slash character is treated as an ordinary character.														
<code>FNM_NOESCAPE</code>	If not set, a backslash character (\) in <i>pattern</i> followed by any other character will match that second character in <i>string</i> . In particular, "\\\" will match a backslash in <i>string</i> .														
	If set, a backslash character will be treated as an ordinary character.														
<code>FNM_PERIOD</code>	If set, a leading period in <i>string</i> will match a period in <i>pattern</i> ; where the location of "leading" is indicated by the value of <code>FNM_PATHNAME</code> :														
	<ul style="list-style-type: none">■ If <code>FNM_PATHNAME</code> is set, a period is "leading" if it is the first character in <i>string</i> or if it immediately follows a slash.■ If <code>FNM_PATHNAME</code> is not set, a period is "leading" only if it is the first character of <i>string</i>.														
	If not set, no special restrictions are placed on matching a period.														
RETURN VALUES	If <i>string</i> matches the pattern specified by <i>pattern</i> , then <code>fnmatch()</code> returns 0. If there is no match, <code>fnmatch()</code> returns <code>FNM_NOMATCH</code> , which is defined in the header <code><fnmatch.h></code> . If an error occurs, <code>fnmatch()</code> returns another non-zero value.														
USAGE	The <code>fnmatch()</code> function has two major uses. It could be used by an application or utility that needs to read a directory and apply a pattern against each entry. The <code>find(1)</code> utility is an example of this. It can also be used by the <code>pax(1)</code> utility to process its <i>pattern</i> operands, or by applications that need to match strings in a similar manner.														

The name `fnmatch()` is intended to imply *filename* match, rather than *pathname* match. The default action of this function is to match filenames, rather than path names, since it gives no special significance to the slash character. With the `FNM_PATHNAME` flag, `fnmatch()` does match path names, but without tilde expansion, parameter expansion, or special treatment for period at the beginning of a filename.

The `fnmatch()` function can be used safely in multithreaded applications, as long as `setlocale(3C)` is not being called to change the locale.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

SEE ALSO `find(1)`, `pax(1)`, `glob(3C)`, `setlocale(3C)`, `wordexp(3C)`, `attributes(5)`, `fnmatch(5)`, `standards(5)`

fopen(3C)

NAME	fopen – open a stream												
SYNOPSIS	<pre>#include <stdio.h> FILE *fopen(const char *filename, const char *mode);</pre>												
DESCRIPTION	<p>The <code>fopen()</code> function opens the file whose pathname is the string pointed to by <i>filename</i>, and associates a stream with it.</p> <p>The argument <i>mode</i> points to a string beginning with one of the following sequences:</p> <table><tr><td>r or rb</td><td>Open file for reading.</td></tr><tr><td>w or wb</td><td>Truncate to zero length or create file for writing.</td></tr><tr><td>a or ab</td><td>Append; open or create file for writing at end-of-file.</td></tr><tr><td>r+ or rb+ or r+b</td><td>Open file for update (reading and writing).</td></tr><tr><td>w+ or wb+ or w+b</td><td>Truncate to zero length or create file for update.</td></tr><tr><td>a+ or ab+ or a+b</td><td>Append; open or create file for update, writing at end-of-file.</td></tr></table> <p>The character <code>b</code> has no effect, but is allowed for ISO C standard conformance (see standards(5)). Opening a file with read mode (<code>r</code> as the first character in the <i>mode</i> argument) fails if the file does not exist or cannot be read.</p> <p>Opening a file with append mode (<code>a</code> as the first character in the <i>mode</i> argument) causes all subsequent writes to the file to be forced to the then current end-of-file, regardless of intervening calls to fseek(3C). If two separate processes open the same file for append, each process may write freely to the file without fear of destroying output being written by the other. The output from the two processes will be intermixed in the file in the order in which it is written.</p> <p>When a file is opened with update mode (<code>+</code> as the second or third character in the <i>mode</i> argument), both input and output may be performed on the associated stream. However, output must not be directly followed by input without an intervening call to fflush(3C) or to a file positioning function (fseek(3C), fsetpos(3C) or rewind(3C)), and input must not be directly followed by output without an intervening call to a file positioning function, unless the input operation encounters end-of-file.</p> <p>When opened, a stream is fully buffered if and only if it can be determined not to refer to an interactive device. The error and end-of-file indicators for the stream are cleared.</p> <p>If <i>mode</i> is <code>w</code>, <code>a</code>, <code>w+</code> or <code>a+</code> and the file did not previously exist, upon successful completion, <code>fopen()</code> function will mark for update the <code>st_atime</code>, <code>st_ctime</code> and <code>st_mtime</code> fields of the file and the <code>st_ctime</code> and <code>st_mtime</code> fields of the parent directory.</p>	r or rb	Open file for reading.	w or wb	Truncate to zero length or create file for writing.	a or ab	Append; open or create file for writing at end-of-file.	r+ or rb+ or r+b	Open file for update (reading and writing).	w+ or wb+ or w+b	Truncate to zero length or create file for update.	a+ or ab+ or a+b	Append; open or create file for update, writing at end-of-file.
r or rb	Open file for reading.												
w or wb	Truncate to zero length or create file for writing.												
a or ab	Append; open or create file for writing at end-of-file.												
r+ or rb+ or r+b	Open file for update (reading and writing).												
w+ or wb+ or w+b	Truncate to zero length or create file for update.												
a+ or ab+ or a+b	Append; open or create file for update, writing at end-of-file.												

If *mode* is *w* or *w+* and the file did previously exist, upon successful completion, `fopen()` will mark for update the `st_ctime` and `st_mtime` fields of the file. The `fopen()` function will allocate a file descriptor as `open(2)` does.

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

RETURN VALUES

Upon successful completion, `fopen()` returns a pointer to the object controlling the stream. Otherwise, a null pointer is returned and `errno` is set to indicate the error.

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

ERRORS

The `fopen()` 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>mode</i> are denied, or the file does not exist and write permission is denied for the parent directory of the file to be created.
EINTR	A signal was caught during the execution of <code>fopen()</code> .
EISDIR	The named file is a directory and <i>mode</i> requires write access.
ELOOP	Too many symbolic links were encountered in resolving <i>path</i> .
EMFILE	There are <code>OPEN_MAX</code> file descriptors currently open in the calling process.
ENAMETOOLONG	The length of the <i>filename</i> 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	A component of <i>filename</i> does not name an existing file or <i>filename</i> is an empty string.
ENOSPC	The directory or file system that would contain the new file cannot be expanded, the file does not exist, and it was to be created.
ENOTDIR	A component of the path prefix is not a directory.
ENXIO	The named file is a character special or block special file, and the device associated with this special file does not exist.
EOVERFLOW	The current value of the file position cannot be represented correctly in an object of type <code>fpos_t</code> .

fopen(3C)

EROFS The named file resides on a read-only file system and *mode* requires write access.

The `fopen()` function may fail if:

EINVAL The value of the *mode* argument is not valid.

EMFILE The number of streams currently open in the calling process is either `FOPEN_MAX` or `STREAM_MAX`.

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

ENOMEM Insufficient storage space is available.

ETXTBSY The file is a pure procedure (shared text) file that is being executed and *mode* requires write access.

USAGE The minimum number of streams that a process can have open at one time is `FOPEN_MAX`. For 32-bit applications, however, the underlying ABIs require that no file descriptor used to access the file underlying an `stdio` `STREAM` have a value greater than 255.

The `fopen()` 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
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `fclose(3C)`, `fdopen(3C)`, `fflush(3C)`, `freopen(3C)`, `fsetpos(3C)`, `rewind(3C)`, `attributes(5)`, `lf64(5)`, `standards(5)`

NAME	fopen, freopen – open a stream
SYNOPSIS	<pre> /usr/ucb/cc[flag ...] file ... #include <stdio.h> FILE *fopen(file, mode); const char *file, *mode; FILE *freopen(file, mode, iop); const char *file, *mode; register FILE *iop; </pre>
DESCRIPTION	<p>The <code>fopen()</code> function opens the file specified by <i>file</i> and associates a stream with it. If the open succeeds, <code>fopen()</code> returns a pointer to be used to identify the stream in subsequent operations. The <i>file</i> argument points to a character string that contains the name of the file to be opened. The <i>mode</i> argument is a character string having one of the following values:</p> <ul style="list-style-type: none"> r open for reading w truncate or create for writing a append: open for writing at end of file, or create for writing r+ open for update (reading and writing) w+ truncate or create for update a+ append; open or create for update at EOF <p>The <code>freopen()</code> function opens the file specified by <i>file</i> and associates the stream pointed to by <i>iop</i> with it. The <i>mode</i> argument is used just as in <code>fopen()</code>. The original stream is closed, regardless of whether the open ultimately succeeds. If the open succeeds, <code>freopen()</code> returns the original value of <i>iop</i>.</p> <p>The <code>freopen()</code> function is typically used to attach the preopened streams associated with <code>stdin</code>, <code>stdout</code>, and <code>stderr</code> to other files.</p> <p>When a file is opened for update, both input and output can be performed on the resulting stream. Output cannot be directly followed by input without an intervening <code>fseek(3C)</code> or <code>rewind(3C)</code>. Input cannot be directly followed by output without an intervening <code>fseek(3C)</code> or <code>rewind(3C)</code>. An input operation that encounters EOF will fail.</p>
RETURN VALUES	The <code>fopen()</code> and <code>freopen()</code> functions return a NULL pointer on failure.
USAGE	The <code>fopen()</code> and <code>freopen()</code> functions have transitional interfaces for 64-bit file offsets. See <code>lf64(5)</code> .
SEE ALSO	<code>open(2)</code> , <code>fclose(3C)</code> , <code>fopen(3C)</code> , <code>freopen(3C)</code> , <code>fseek(3C)</code> , <code>malloc(3C)</code> , <code>rewind(3C)</code> , <code>lf64(5)</code>

fopen(3UCB)

NOTES Use of these functions should be restricted to applications written on BSD platforms. Use of these functions with any of the system libraries or in multithreaded applications is unsupported.

To support the same number of open files as the system, `fopen()` must allocate additional memory for data structures using `malloc(3C)` after 64 files have been opened. This confuses some programs that use their own memory allocators.

The `fopen()` and `freopen()` functions differ from the standard I/O functions `fopen(3C)` and `freopen(3C)`. The standard I/O functions distinguish binary from text files with an additional use of 'b' as part of the *mode*, enabling portability of `fopen(3C)` and `freopen(3C)` beyond SunOS 4.x systems.

NAME	fpgetround, fpsetround, fpgetmask, fpsetmask, fpgetsticky, fpsetsticky – IEEE floating-point environment control
SYNOPSIS	<pre>#include <ieeefp.h> fp_rnd fpgetround(void); fp_rnd fpsetround(fp_rnd rnd_dir); fp_except fpgetmask(void); fp_except fpsetmask(fp_except mask); fp_except fpgetsticky(void); fp_except fpsetsticky(fp_except sticky);</pre>
DESCRIPTION	<p>There are five floating-point exceptions:</p> <ul style="list-style-type: none"> ■ divide-by-zero, ■ overflow, ■ underflow, ■ imprecise (inexact) result, and ■ invalid operation. <p>When a floating-point exception occurs, the corresponding sticky bit is set (1), and if the mask bit is enabled (1), the trap takes place. These routines let the user change the behavior on occurrence of any of these exceptions, as well as change the rounding mode for floating-point operations.</p> <p>The <i>mask</i> argument is formed by the logical OR operation of the following floating-point exception masks:</p> <pre>FP_X_INV /* invalid operation exception */ FP_X_OFL /* overflow exception */ FP_X_UFL /* underflow exception */ FP_X_DZ /* divide-by-zero exception */ FP_X_IMP /* imprecise (loss of precision) */</pre> <p>The following floating-point rounding modes are passed to <code>fpsetround</code> and returned by <code>fpgetround()</code>.</p> <pre>FP_RN /* round to nearest representative number */ FP_RP /* round to plus infinity */ FP_RM /* round to minus infinity */ FP_RZ /* round to zero (truncate) */</pre> <p>The default environment is rounding mode set to nearest (<code>FP_RN</code>) and all traps disabled.</p> <p>The <code>fpsetsticky()</code> function modifies all sticky flags. The <code>fpsetmask()</code> function changes all mask bits. The <code>fpsetmask()</code> function clears the sticky bit corresponding to any exception being enabled.</p>
RETURN VALUES	The <code>fpgetround()</code> function returns the current rounding mode.

fpgetround(3C)

The `fpsetround()` function sets the rounding mode and returns the previous rounding mode.

The `fpgetmask()` function returns the current exception masks.

The `fpsetmask()` function sets the exception masks and returns the previous setting.

The `fpgetsticky()` function returns the current exception sticky flags.

The `fpsetsticky()` function sets (clears) the exception sticky flags and returns the previous setting.

USAGE The C programming language requires truncation (round to zero) for floating point to integral conversions. The current rounding mode has no effect on these conversions.

The sticky bit must be cleared to recover from the trap and proceed. If the sticky bit is not cleared before the next trap occurs, a wrong exception type may be signaled.

Individual bits may be examined using the constants defined in `<ieeefp.h>`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [isnand\(3C\)](#), [attributes\(5\)](#)

NAME	fputc, putc, putc_unlocked, putchar, putchar_unlocked, putw – put a byte on a stream
SYNOPSIS	<pre>#include <stdio.h> int fputc(int c, FILE *stream); int putc(int c, FILE *stream); int putc_unlocked(int c, FILE *stream); int putchar(int c); int putchar_unlocked(int c); int putw(int w, FILE *stream);</pre>
DESCRIPTION	<p>The <code>fputc()</code> function writes the byte specified by <code>c</code> (converted to an unsigned char) to the output stream pointed to by <code>stream</code>, at the position indicated by the associated file-position indicator for the stream (if defined), and advances the indicator appropriately. If the file cannot support positioning requests, or if the stream was opened with append mode, the byte is appended to the output stream.</p> <p>The <code>st_ctime</code> and <code>st_mtime</code> fields of the file will be marked for update between the successful execution of <code>fputc()</code> and the next successful completion of a call to <code>fflush(3C)</code> or <code>fclose(3C)</code> on the same stream or a call to <code>exit(3C)</code> or <code>abort(3C)</code>.</p> <p>The <code>putc()</code> routine behaves like <code>fputc()</code>, except that it is implemented as a macro. It runs faster than <code>fputc()</code>, but it takes up more space per invocation and its name cannot be passed as an argument to a function call.</p> <p>The call <code>putchar(c)</code> is equivalent to <code>putc(c, stdout)</code>. The <code>putchar()</code> routine is implemented as a macro.</p> <p>The <code>putc_unlocked()</code> and <code>putchar_unlocked()</code> routines are variants of <code>putc()</code> and <code>putchar()</code>, respectively, that do not lock the stream. It is the caller's responsibility to acquire the stream lock before calling these routines and releasing the lock afterwards; see <code>flockfile(3C)</code> and <code>stdio(3C)</code>. These routines are implemented as macros.</p> <p>The <code>putw()</code> function writes the word (that is, type <code>int</code>) <code>w</code> to the output <code>stream</code> (at the position at which the file offset, if defined, is pointing). The size of a word is the size of a type <code>int</code> and varies from machine to machine. The <code>putw()</code> function neither assumes nor causes special alignment in the file.</p> <p>The <code>st_ctime</code> and <code>st_mtime</code> fields of the file will be marked for update between the successful execution of <code>putw()</code> and the next successful completion of a call to <code>fflush(3C)</code> or <code>fclose(3C)</code> on the same stream or a call to <code>exit(3C)</code> or <code>abort(3C)</code>.</p>
RETURN VALUES	Upon successful completion, <code>fputc()</code> , <code>putc()</code> , <code>putc_unlocked()</code> , <code>putchar()</code> , and <code>putchar_unlocked()</code> return the value that was written. Otherwise, these functions return EOF, the error indicator for the stream is set, and <code>errno</code> is set to indicate the error.

fputc(3C)

Upon successful completion, `putw()` returns 0. Otherwise, it returns a non-zero value, sets the error indicator for the associated *stream*, and sets `errno` to indicate the error.

An unsuccessful completion will occur, for example, if the file associated with *stream* is not open for writing or if the output file cannot grow.

ERRORS

The `fputc()`, `putc()`, `putc_unlocked()`, `putchar()`, `putchar_unlocked()`, and `putw()` functions will fail if either the *stream* is unbuffered or the *stream*'s buffer needs to be flushed, and:

EAGAIN	The <code>O_NONBLOCK</code> flag is set for the file descriptor underlying <i>stream</i> and the process would be delayed in the write operation.
EBADF	The file descriptor underlying <i>stream</i> is not a valid file descriptor open for writing.
EFBIG	An attempt was made to write to a file that exceeds the maximum file size or the process' file size limit.
EFBIG	The file is a regular file and an attempt was made to write at or beyond the offset maximum.
EINTR	The write operation was terminated due to the receipt of a signal, and no data was transferred.
EIO	A physical I/O error has occurred, or the process is a member of a background process group attempting to write to its controlling terminal, <code>TOSTOP</code> is set, the process is neither ignoring nor blocking <code>SIGTTOU</code> and the process group of the process is orphaned. This error may also be returned under implementation-dependent conditions.
ENOSPC	There was no free space remaining on the device containing the file.
EPIPE	An attempt is made to write to a pipe or FIFO that is not open for reading by any process. A <code>SIGPIPE</code> signal will also be sent to the calling thread.

The `fputc()`, `putc()`, `putc_unlocked()`, `putchar()`, `putchar_unlocked()`, and `putw()` functions may fail if:

ENOMEM	Insufficient storage space is available.
ENXIO	A request was made of a non-existent device, or the request was outside the capabilities of the device.

USAGE

Functions exist for the `putc()`, `putc_unlocked()`, `putchar()`, and `putchar_unlocked()` macros. To get the function form, the macro name must be undefined (for example, `#undef putc`).

When the macro forms are used, `putc()` and `putc_unlocked()` evaluate the *stream* argument more than once. In particular, `putc(c, *f++)`; does not work sensibly. The `fputc()` function should be used instead when evaluating the *stream* argument has side effects.

Because of possible differences in word length and byte ordering, files written using `putw()` are implementation-dependent, and possibly cannot be read using `getw(3C)` by a different application or by the same application running in a different environment.

The `putw()` function is inherently byte stream oriented and is not tenable in the context of either multibyte character streams or wide-character streams. Application programmers are encouraged to use one of the character-based output functions instead.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	<code>fputc()</code> , <code>putc()</code> , <code>putc_unlocked()</code> , <code>putchar()</code> , and <code>putchar_unlocked()</code> are Standard.
MT-Level	See NOTES below.

SEE ALSO `getrlimit(2)`, `ulimit(2)`, `write(2)`, `intro(3)`, `abort(3C)`, `exit(3C)`, `fclose(3C)`, `ferror(3C)`, `fflush(3C)`, `flockfile(3C)`, `fopen(3UCB)`, `printf(3C)`, `putc(3C)`, `puts(3C)`, `setbuf(3C)`, `stdio(3C)`, `attributes(5)`, `standards(5)`

NOTES The `fputc()`, `putc()`, `putchar()`, and `putw()` routines are MT-Safe in multithreaded applications. The `putc_unlocked()` and `putchar_unlocked()` routines are unsafe in multithreaded applications.

fputwc(3C)

NAME	fputwc, putwc, putwchar – put wide-character code on a stream
SYNOPSIS	<pre>#include <stdio.h> #include <wchar.h> wint_t fputwc(wchar_t <i>wc</i>, FILE*<i>stream</i>); wint_t putwc(wchar_t <i>wc</i>, FILE*<i>stream</i>); #include <wchar.h> wint_t putwchar(wchar_t <i>wc</i>);</pre>
DESCRIPTION	<p>The <code>fputwc()</code> function writes the character corresponding to the wide-character code <i>wc</i> to the output stream pointed to by <i>stream</i>, at the position indicated by the associated file-position indicator for the stream (if defined), and advances the indicator appropriately. If the file cannot support positioning requests, or if the stream was opened with append mode, the character is appended to the output stream. If an error occurs while writing the character, the shift state of the output file is left in an undefined state.</p> <p>The <code>st_ctime</code> and <code>st_mtime</code> fields of the file will be marked for update between the successful execution of <code>fputwc()</code> and the next successful completion of a call to <code>fflush(3C)</code> or <code>fclose(3C)</code> on the same stream or a call to <code>exit(2)</code> or <code>abort(3C)</code>.</p> <p>The <code>putwc()</code> function is equivalent to <code>fputwc()</code>, except that it is implemented as a macro.</p> <p>The call <code>putwchar(<i>wc</i>)</code> is equivalent to <code>putwc(<i>wc</i>, stdout)</code>. The <code>putwchar()</code> routine is implemented as a macro.</p>
RETURN VALUES	Upon successful completion, <code>fputwc()</code> , <code>putwc()</code> , and <code>putwchar()</code> return <i>wc</i> . Otherwise, they return <code>WEOF</code> , the error indicator for the stream is set, and <code>errno</code> is set to indicate the error.
ERRORS	The <code>fputwc()</code> , <code>putwc()</code> , and <code>putwchar()</code> functions will fail if either the stream is unbuffered or data in the <i>stream</i> 's buffer needs to be written, and:
EAGAIN	The <code>O_NONBLOCK</code> flag is set for the file descriptor underlying <i>stream</i> and the process would be delayed in the write operation.
EBADF	The file descriptor underlying <i>stream</i> is not a valid file descriptor open for writing.
EFBIG	An attempt was made to write to a file that exceeds the maximum file size or the process's file size limit; or the file is a regular file and an attempt was made to write at or beyond the offset maximum associated with the corresponding stream.
EINTR	The write operation was terminated due to the receipt of a signal, and no data was transferred.
EIO	A physical I/O error has occurred, or the process is a member of a background process group attempting to write to its controlling

terminal, TOSTOP is set, the process is neither ignoring nor blocking SIGTTOU, and the process group of the process is orphaned.

ENOSPC There was no free space remaining on the device containing the file.

EPIPE An attempt is made to write to a pipe or FIFO that is not open for reading by any process. A SIGPIPE signal will also be sent to the calling thread.

The fputc(), putc(), and putchar() functions may fail if:

ENOMEM Insufficient storage space is available.

ENXIO A request was made of a non-existent device, or the request was outside the capabilities of the device.

EILSEQ The wide-character code *wc* does not correspond to a valid character.

USAGE Functions exist for the putc() and putchar() macros. To get the function form, the macro name must be undefined (for example, #undef putc).

When the macro form is used, putc() evaluates the *stream* argument more than once. In particular, putc(wc, *f++) does not work sensibly. The fputc() function should be used instead when evaluating the *stream* argument has side effects.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO exit(2), ulimit(2), abort(3C), fclose(3C), ferrord(3C), fflush(3C), fopen(3C), setbuf(3C), attributes(5), standards(5)

fputws(3C)

NAME	fputws – put wide character string on a stream						
SYNOPSIS	<pre>#include <stdio.h> #include <wchar.h> int fputws(const wchar_t *restrict s, FILE *restrict stream);</pre>						
DESCRIPTION	<p>The <code>fputws()</code> function writes a character string corresponding to the (null-terminated) wide character string pointed to by <code>ws</code> to the stream pointed to by <code>stream</code>. No character corresponding to the terminating null wide-character code is written, nor is a <code>NEWLINE</code> character appended.</p> <p>The <code>st_ctime</code> and <code>st_mtime</code> fields of the file will be marked for update between the successful execution of <code>fputws()</code> and the next successful completion of a call to <code>fflush(3C)</code> or <code>fclose(3C)</code> on the same stream or a call to <code>exit(2)</code> or <code>abort(3C)</code>.</p>						
RETURN VALUES	Upon successful completion, <code>fputws()</code> returns a non-negative value. Otherwise, it returns <code>-1</code> , sets an error indicator for the stream, and sets <code>errno</code> to indicate the error.						
ERRORS	Refer to <code>fputwc(3C)</code> .						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>exit(2)</code> , <code>abort(3C)</code> , <code>fclose(3C)</code> , <code>fflush(3C)</code> , <code>fopen(3C)</code> , <code>fputwc(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	fread – binary input
SYNOPSIS	<pre>#include <stdio.h> size_t fread(void *ptr, size_t size, size_t nitems, FILE *stream);</pre>
DESCRIPTION	<p>The <code>fread()</code> function reads into the array pointed to by <i>ptr</i> up to <i>nitems</i> elements whose size is specified by <i>size</i> in bytes, from the stream pointed to by <i>stream</i>. For each object, <i>size</i> calls are made to the <code>fgetc(3C)</code> function and the results stored, in the order read, in an array of unsigned char exactly overlaying the object. The file-position indicator for the stream (if defined) is advanced by the number of bytes successfully read. If an error occurs, the resulting value of the file-position indicator for the stream is unspecified. If a partial element is read, its value is unspecified.</p> <p>The <code>fread()</code> function may mark the <code>st_atime</code> field of the file associated with <i>stream</i> for update. The <code>st_atime</code> field will be marked for update by the first successful execution of <code>fgetc(3C)</code>, <code>fgets(3C)</code>, <code>fgetwc(3C)</code>, <code>fgetws(3C)</code>, <code>fread()</code>, <code>fscanf(3C)</code>, <code>getc(3C)</code>, <code>getchar(3C)</code>, <code>gets(3C)</code>, or <code>scanf(3C)</code> using <i>stream</i> that returns data not supplied by a prior call to <code>ungetc(3C)</code> or <code>ungetwc(3C)</code>.</p>
RETURN VALUES	<p>Upon successful completion, <code>fread()</code> returns the number of elements successfully read, which is less than <i>nitems</i> only if a read error or end-of-file is encountered. If <i>size</i> or <i>nitems</i> is 0, <code>fread()</code> returns 0 and the contents of the array and the state of the stream remain unchanged. Otherwise, if a read error occurs, the error indicator for the stream is set and <code>errno</code> is set to indicate the error.</p>
ERRORS	Refer to <code>fgetc(3C)</code> .
EXAMPLES	<p>EXAMPLE 1 Reading from a Stream</p> <p>The following example reads a single element from the <i>fp</i> stream into the array pointed to by <i>buf</i>.</p> <pre>#include <stdio.h> ... size_t bytes_read; char buf[100]; FILE *fp; ... bytes_read = fread(buf, sizeof(buf), 1, fp); ...</pre>
USAGE	<p>The <code>ferror()</code> or <code>feof()</code> functions must be used to distinguish between an error condition and end-of-file condition. See <code>ferror(3C)</code>.</p> <p>Because of possible differences in element length and byte ordering, files written using <code>fwrite(3C)</code> are application-dependent, and possibly cannot be read using <code>fread()</code> by a different application or by the same application on a different processor.</p>

fread(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `read(2)`, `fclose(3C)`, `ferror(3C)`, `fopen(3C)`, `getc(3C)`, `gets(3C)`, `printf(3C)`, `putc(3C)`, `puts(3C)`, `attributes(5)`, `standards(5)`

NAME	freopen – open a stream				
SYNOPSIS	<pre>#include <stdio.h> FILE *freopen(const char *filename, const char *mode, FILE *stream);</pre>				
DESCRIPTION	<p>The <code>freopen()</code> function first attempts to flush the stream and close any file descriptor associated with <i>stream</i>. Failure to flush or close the file successfully is ignored. The error and end-of-file indicators for the stream are cleared.</p> <p>The <code>freopen()</code> function opens the file whose pathname is the string pointed to by <i>filename</i> and associates the stream pointed to by <i>stream</i> with it. The <i>mode</i> argument is used just as in <code>fopen(3C)</code>.</p> <p>If <i>filename</i> is a null pointer and the application conforms to SUSv3 (see <code>standards(5)</code>), the <code>freopen()</code> function attempts to change the mode of the stream to that specified by <i>mode</i>, as though the name of the file currently associated with the <i>stream</i> had been used. The following changes of mode are permitted, depending upon the access mode of the file descriptor underlying the stream:</p> <ul style="list-style-type: none"> ■ When <code>+</code> is specified, the file descriptor mode must be <code>O_RDWR</code>. ■ When <code>r</code> is specified, the file descriptor mode must be <code>O_RDONLY</code> or <code>O_RDWR</code>. ■ When <code>a</code> or <code>w</code> is specified, the file descriptor mode must be <code>O_WRONLY</code> or <code>O_RDWR</code>. <p>If the <i>filename</i> is a null pointer and the application does not conform to SUSv3, <code>freopen()</code> returns a null pointer.</p> <p>The original stream is closed regardless of whether the subsequent open succeeds.</p> <p>After a successful call to the <code>freopen()</code> function, the orientation of the stream is cleared, the encoding rule is cleared, and the associated <code>mbstate_t</code> object is set to describe an initial conversion state.</p> <p>The largest value that can be represented correctly in an object of type <code>off_t</code> will be established as the offset maximum in the open file description.</p>				
RETURN VALUES	Upon successful completion, <code>freopen()</code> returns the value of <i>stream</i> . Otherwise, a null pointer is returned and <code>errno</code> is set to indicate the error.				
ERRORS	<p>The <code>freopen()</code> function will fail if:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">EACCES</td> <td>Search permission is denied on a component of the path prefix, or the file exists and the permissions specified by <i>mode</i> are denied, or the file does not exist and write permission is denied for the parent directory of the file to be created.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">EBADF</td> <td>The application conforms to SUSv3, the <i>filename</i> argument is a null pointer, and either the underlying file descriptor is not valid or the mode specified when the underlying file descriptor was opened does not support the file access modes requested by the <i>mode</i> argument.</td> </tr> </table>	EACCES	Search permission is denied on a component of the path prefix, or the file exists and the permissions specified by <i>mode</i> are denied, or the file does not exist and write permission is denied for the parent directory of the file to be created.	EBADF	The application conforms to SUSv3, the <i>filename</i> argument is a null pointer, and either the underlying file descriptor is not valid or the mode specified when the underlying file descriptor was opened does not support the file access modes requested by the <i>mode</i> argument.
EACCES	Search permission is denied on a component of the path prefix, or the file exists and the permissions specified by <i>mode</i> are denied, or the file does not exist and write permission is denied for the parent directory of the file to be created.				
EBADF	The application conforms to SUSv3, the <i>filename</i> argument is a null pointer, and either the underlying file descriptor is not valid or the mode specified when the underlying file descriptor was opened does not support the file access modes requested by the <i>mode</i> argument.				

freopen(3C)

EFAULT	The application does not conform to SUSv3 and the <i>filename</i> argument is a null pointer.
EINTR	A signal was caught during <code>freopen()</code> .
EISDIR	The named file is a directory and <i>mode</i> requires write access.
ELOOP	Too many symbolic links were encountered in resolving <i>path</i> .
EMFILE	There are {OPEN_MAX} file descriptors currently open in the calling process.
ENAMETOOLONG	The length of the <i>filename</i> exceeds {PATH_MAX} or a pathname component is longer than {NAME_MAX}.
ENFILE	The maximum allowable number of files is currently open in the system.
ENOENT	A component of <i>filename</i> does not name an existing file or <i>filename</i> is an empty string.
ENOSPC	The directory or file system that would contain the new file cannot be expanded, the file does not exist, and it was to be created.
ENOTDIR	A component of the path prefix is not a directory.
ENXIO	The named file is a character special or block special file, and the device associated with this special file does not exist.
E_OVERFLOW	The current value of the file position cannot be represented correctly in an object of type <code>off_t</code> .
EROFS	The named file resides on a read-only file system and <i>mode</i> requires write access.

The `freopen()` function may fail if:

EINVAL	The value of the <i>mode</i> argument is not valid.
ENAMETOOLONG	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds {PATH_MAX}.
ENOMEM	Insufficient storage space is available.
ENXIO	A request was made of a non-existent device, or the request was outside the capabilities of the device.
ETXTBSY	The file is a pure procedure (shared text) file that is being executed and <i>mode</i> requires write access.

USAGE The `freopen()` function is typically used to attach the preopened *streams* associated with `stdin`, `stdout` and `stderr` to other files. By default `stderr` is unbuffered, but the use of `freopen()` will cause it to become buffered or line-buffered.

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

freopen(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `fclose(3C)`, `fdopen(3C)`, `fopen(3C)`, `stdio(3C)`, `attributes(5)`, `lf64(5)`, `standards(5)`

fseek(3C)

NAME	fseek, fseeko – reposition a file-position indicator in a stream						
SYNOPSIS	<pre>#include <stdio.h> int fseek(FILE *stream, long offset, int whence); int fseeko(FILE *stream, off_t offset, int whence);</pre>						
DESCRIPTION	<p>The <code>fseek()</code> function sets the file-position indicator for the stream pointed to by <code>stream</code>. The <code>fseeko()</code> function is identical to <code>fseek()</code> except for the type of <code>offset</code>.</p> <p>The new position, measured in bytes from the beginning of the file, is obtained by adding <code>offset</code> to the position specified by <code>whence</code>, whose values are defined in <code><stdio.h></code> as follows:</p> <table><tr><td>SEEK_SET</td><td>Set position equal to <code>offset</code> bytes.</td></tr><tr><td>SEEK_CUR</td><td>Set position to current location plus <code>offset</code>.</td></tr><tr><td>SEEK_END</td><td>Set position to EOF plus <code>offset</code>.</td></tr></table> <p>If the stream is to be used with wide character input/output functions, <code>offset</code> must either be 0 or a value returned by an earlier call to <code>ftell(3C)</code> on the same stream and <code>whence</code> must be <code>SEEK_SET</code>.</p> <p>A successful call to <code>fseek()</code> clears the end-of-file indicator for the stream and undoes any effects of <code>ungetc(3C)</code> and <code>ungetwc(3C)</code> on the same stream. After an <code>fseek()</code> call, the next operation on an update stream may be either input or output.</p> <p>If the most recent operation, other than <code>ftell(3C)</code>, on a given stream is <code>fflush(3C)</code>, the file offset in the underlying open file description will be adjusted to reflect the location specified by <code>fseek()</code>.</p> <p>The <code>fseek()</code> function allows the file-position indicator to be set beyond the end of existing data in the file. If data is later written at this point, subsequent reads of data in the gap will return bytes with the value 0 until data is actually written into the gap.</p> <p>The value of the file offset returned by <code>fseek()</code> on devices which are incapable of seeking is undefined.</p> <p>If the stream is writable and buffered data had not been written to the underlying file, <code>fseek()</code> will cause the unwritten data to be written to the file and mark the <code>st_ctime</code> and <code>st_mtime</code> fields of the file for update.</p>	SEEK_SET	Set position equal to <code>offset</code> bytes.	SEEK_CUR	Set position to current location plus <code>offset</code> .	SEEK_END	Set position to EOF plus <code>offset</code> .
SEEK_SET	Set position equal to <code>offset</code> bytes.						
SEEK_CUR	Set position to current location plus <code>offset</code> .						
SEEK_END	Set position to EOF plus <code>offset</code> .						
RETURN VALUES	The <code>fseek()</code> and <code>fseeko()</code> functions return 0 on success; otherwise, they returned -1 and set <code>errno</code> to indicate the error.						
ERRORS	<p>The <code>fseek()</code> and <code>fseeko()</code> functions will fail if, either the <code>stream</code> is unbuffered or the <code>stream</code>'s buffer needed to be flushed, and the call to <code>fseek()</code> or <code>fseeko()</code> causes an underlying <code>lseek(2)</code> or <code>write(2)</code> to be invoked:</p> <table><tr><td>EAGAIN</td><td>The <code>O_NONBLOCK</code> flag is set for the file descriptor and the process would be delayed in the write operation.</td></tr></table>	EAGAIN	The <code>O_NONBLOCK</code> flag is set for the file descriptor and the process would be delayed in the write operation.				
EAGAIN	The <code>O_NONBLOCK</code> flag is set for the file descriptor and the process would be delayed in the write operation.						

EBADF	The file descriptor underlying the stream file is not open for writing or the stream's buffer needed to be flushed and the file is not open.
EFBIG	An attempt was made to write a file that exceeds the maximum file size or the process's file size limit, or the file is a regular file and an attempt was made to write at or beyond the offset maximum associated with the corresponding stream.
EINTR	The write operation was terminated due to the receipt of a signal, and no data was transferred.
EINVAL	The <i>whence</i> argument is invalid. The resulting file-position indicator would be set to a negative value.
EIO	A physical I/O error has occurred; or the process is a member of a background process group attempting to perform a <code>write(2)</code> operation to its controlling terminal, <code>TOSTOP</code> is set, the process is neither ignoring nor blocking <code>SIGTTOU</code> , and the process group of the process is orphaned.
ENOSPC	There was no free space remaining on the device containing the file.
ENXIO	A request was made of a non-existent device, or the request was outside the capabilities of the device.
EPIPE	The file descriptor underlying <i>stream</i> is associated with a pipe or FIFO.
EPIPE	An attempt was made to write to a pipe or FIFO that is not open for reading by any process. A <code>SIGPIPE</code> signal will also be sent to the calling thread.

The `fseek()` function will fail if:

E_OVERFLOW	The resulting file offset would be a value which cannot be represented correctly in an object of type <code>long</code> .
------------	---

The `fseeko()` function will fail if:

E_OVERFLOW	The resulting file offset would be a value which cannot be represented correctly in an object of type <code>off_t</code> .
------------	--

USAGE Although on the UNIX system an offset returned by `ftell()` or `ftello()` (see [ftell\(3C\)](#)) is measured in bytes, and it is permissible to seek to positions relative to that offset, portability to non-UNIX systems requires that an offset be used by `fseek()` directly. Arithmetic may not meaningfully be performed on such an offset, which is not necessarily measured in bytes.

The `fseeko()` function has a transitional interface for 64-bit file offsets. See [1f64\(5\)](#).

fseek(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `getrlimit(2)`, `ulimit(2)`, `fopen(3UCB)`, `ftell(3C)`, `rewind(3C)`, `ungetc(3C)`, `ungetwc(3C)`, `attributes(5)`, `lf64(5)`, `standards(5)`

NAME fsetpos – reposition a file pointer in a stream

SYNOPSIS

```
#include <stdio.h>

int fsetpos(FILE *stream, const fpos_t *pos);
```

DESCRIPTION The `fsetpos()` function sets the file position indicator for the stream pointed to by `stream` according to the value of the object pointed to by `pos`, which must be a value obtained from an earlier call to `fgetpos(3C)` on the same stream.

A successful call to `fsetpos()` function clears the end-of-file indicator for the stream and undoes any effects of `ungetc(3C)` on the same stream. After an `fsetpos()` call, the next operation on an update stream may be either input or output.

RETURN VALUES The `fsetpos()` function returns 0 if it succeeds; otherwise it returns a non-zero value and sets `errno` to indicate the error.

ERRORS The `fsetpos()` function may fail if:

EBADF	The file descriptor underlying <code>stream</code> is not valid.
ESPIPE	The file descriptor underlying <code>stream</code> is associated with a pipe, a FIFO, or a socket.

USAGE The `fsetpos()` 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
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `lseek(2)`, `fgetpos(3C)`, `fopen(3C)`, `fseek(3C)`, `ftell(3C)`, `rewind(3C)`, `ungetc(3C)`, `attributes(5)`, `lf64(5)`, `standards(5)`

fsync(3C)

NAME	fsync – synchronize changes to a file										
SYNOPSIS	<pre>#include <unistd.h> int fsync (int <i>fildev</i>);</pre>										
DESCRIPTION	<p>The <code>fsync()</code> function moves all modified data and attributes of the file descriptor <i>fildev</i> to a storage device. When <code>fsync()</code> returns, all in-memory modified copies of buffers associated with <i>fildev</i> have been written to the physical medium. The <code>fsync()</code> function is different from <code>sync()</code>, which schedules disk I/O for all files but returns before the I/O completes. The <code>fsync()</code> function forces all outstanding data operations to synchronized file integrity completion (see <code>fcntl.h(3HEAD)</code> definition of <code>O_SYNC</code>.)</p> <p>The <code>fsync()</code> function forces all currently queued I/O operations associated with the file indicated by the file descriptor <i>fildev</i> to the synchronized I/O completion state. All I/O operations are completed as defined for synchronized I/O file integrity completion.</p>										
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error. If the <code>fsync()</code> function fails, outstanding I/O operations are not guaranteed to have been completed.										
ERRORS	<p>The <code>fsync()</code> function will fail if:</p> <table><tr><td>EBADF</td><td>The <i>fildev</i> argument is not a valid file descriptor.</td></tr><tr><td>EINTR</td><td>A signal was caught during execution of the <code>fsync()</code> function.</td></tr><tr><td>EIO</td><td>An I/O error occurred while reading from or writing to the file system.</td></tr><tr><td>ENOSPC</td><td>There was no free space remaining on the device containing the file.</td></tr><tr><td>ETIMEDOUT</td><td>Remote connection timed out. This occurs when the file is on an NFS file system mounted with the <i>soft</i> option. See <code>mount_nfs(1M)</code>.</td></tr></table> <p>In the event that any of the queued I/O operations fail, <code>fsync()</code> returns the error conditions defined for <code>read(2)</code> and <code>write(2)</code>.</p>	EBADF	The <i>fildev</i> argument is not a valid file descriptor.	EINTR	A signal was caught during execution of the <code>fsync()</code> function.	EIO	An I/O error occurred while reading from or writing to the file system.	ENOSPC	There was no free space remaining on the device containing the file.	ETIMEDOUT	Remote connection timed out. This occurs when the file is on an NFS file system mounted with the <i>soft</i> option. See <code>mount_nfs(1M)</code> .
EBADF	The <i>fildev</i> argument is not a valid file descriptor.										
EINTR	A signal was caught during execution of the <code>fsync()</code> function.										
EIO	An I/O error occurred while reading from or writing to the file system.										
ENOSPC	There was no free space remaining on the device containing the file.										
ETIMEDOUT	Remote connection timed out. This occurs when the file is on an NFS file system mounted with the <i>soft</i> option. See <code>mount_nfs(1M)</code> .										
USAGE	<p>The <code>fsync()</code> function should be used by applications that require that a file be in a known state. For example, an application that contains a simple transaction facility might use <code>fsync()</code> to ensure that all changes to a file or files caused by a given transaction were recorded on a storage medium.</p> <p>The manner in which the data reach the physical medium depends on both implementation and hardware. The <code>fsync()</code> function returns when notified by the device driver that the write has taken place.</p>										

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

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

SEE ALSO `mount_nfs(1M)`, `read(2)`, `sync(2)`, `write(2)`, `fcntl.h(3HEAD)`, `fdatasync(3RT)`, `attributes(5)`, `standards(5)`

ftell(3C)

NAME	ftell, ftello – return a file offset in a stream						
SYNOPSIS	<pre>#include <stdio.h> long ftell(FILE *<i>stream</i>) ; off_t ftello(FILE *<i>stream</i>) ;</pre>						
DESCRIPTION	The <code>ftell()</code> function obtains the current value of the file-position indicator for the stream pointed to by <i>stream</i> . The <code>ftello()</code> function is identical to <code>ftell()</code> except for the return type.						
RETURN VALUES	Upon successful completion, the <code>ftell()</code> and <code>ftello()</code> functions return the current value of the file-position indicator for the stream measured in bytes from the beginning of the file. Otherwise, they return <code>-1</code> and sets <code>errno</code> to indicate the error.						
ERRORS	<p>The <code>ftell()</code> and <code>ftello()</code> functions will fail if:</p> <p><code>EBADF</code> The file descriptor underlying <i>stream</i> is not an open file descriptor.</p> <p><code>ESPIPE</code> The file descriptor underlying <i>stream</i> is associated with a pipe, a FIFO, or a socket.</p> <p>The <code>ftell()</code> function will fail if:</p> <p><code>Eoverflow</code> The current file offset cannot be represented correctly in an object of type <code>long</code>.</p> <p>The <code>ftello()</code> function will fail if:</p> <p><code>Eoverflow</code> The current file offset cannot be represented correctly in an object of type <code>off_t</code>.</p>						
USAGE	The <code>ftello()</code> function has a transitional interface for 64-bit file offsets. See <code>lf64(5)</code> .						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>lseek(2)</code> , <code>fopen(3C)</code> , <code>fseek(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code> , <code>lf64(5)</code>						

NAME	ftime – get date and time				
SYNOPSIS	<pre>#include <sys/timeb.h> int ftime(struct timeb *tp);</pre>				
DESCRIPTION	<p>The <code>ftime()</code> function sets the <code>time</code> and <code>millitm</code> members of the <code>timeb</code> structure pointed to by <code>tp</code>. The structure is defined in <code><sys/timeb.h></code> and contains the following members:</p> <pre>time_t time; unsigned short millitm; short timezone; short dstflag;</pre> <p>The <code>time</code> and <code>millitm</code> members contain the seconds and milliseconds portions, respectively, of the current time in seconds since 00:00:00 UTC (Coordinated Universal Time), January 1, 1970.</p> <p>The <code>timezone</code> member contains the local time zone. The <code>dstflag</code> member contains a flag that, if non-zero, indicates that Daylight Saving time applies locally during the appropriate part of the year.</p> <p>The contents of the <code>timezone</code> and <code>dstflag</code> members of <code>tp</code> after a call to <code>ftime()</code> are unspecified.</p>				
RETURN VALUES	Upon successful completion, the <code>ftime()</code> function returns 0. Otherwise <code>-1</code> is returned.				
ERRORS	No errors are defined.				
USAGE	<p>For portability to implementations conforming to earlier versions of this document, <code>time(2)</code> is preferred over this function.</p> <p>The millisecond value usually has a granularity greater than one due to the resolution of the system clock. Depending on any granularity (particularly a granularity of one) renders code non-portable.</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Standard				
SEE ALSO	<code>date(1)</code> , <code>time(2)</code> , <code>ctime(3C)</code> , <code>gettimeofday(3C)</code> , <code>timezone(4)</code> , <code>attributes(5)</code> , <code>standards(5)</code>				

ftok(3C)

NAME	ftok – generate an IPC key												
SYNOPSIS	<pre>#include <sys/ipc.h> key_t ftok(const char *path, int id);</pre>												
DESCRIPTION	<p>The <code>ftok()</code> function returns a key based on <i>path</i> and <i>id</i> that is usable in subsequent calls to <code>msgget(2)</code>, <code>semget(2)</code> and <code>shmget(2)</code>. The <i>path</i> argument must be the pathname of an existing file that the process is able to <code>stat(2)</code>.</p> <p>The <code>ftok()</code> function will return the same key value for all paths that name the same file, when called with the same <i>id</i> value, and will return different key values when called with different <i>id</i> values.</p> <p>If the file named by <i>path</i> is removed while still referred to by a key, a call to <code>ftok()</code> with the same <i>path</i> and <i>id</i> returns an error. If the same file is recreated, then a call to <code>ftok()</code> with the same <i>path</i> and <i>id</i> is likely to return a different key.</p> <p>Only the low order 8-bits of <i>id</i> are significant. The behavior of <code>ftok()</code> is unspecified if these bits are 0.</p>												
RETURN VALUES	Upon successful completion, <code>ftok()</code> returns a key. Otherwise, <code>ftok()</code> returns <code>(key_t)-1</code> and sets <code>errno</code> to indicate the error.												
ERRORS	<p>The <code>ftok()</code> function will fail if:</p> <table><tr><td>EACCES</td><td>Search permission is denied for a component of the path prefix.</td></tr><tr><td>ELOOP</td><td>Too many symbolic links were encountered in resolving <i>path</i>.</td></tr><tr><td>ENAMETOOLONG</td><td>The length of the <i>path</i> argument exceeds <code>{PATH_MAX}</code> or a pathname component is longer than <code>{NAME_MAX}</code>.</td></tr><tr><td>ENOENT</td><td>A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.</td></tr><tr><td>ENOTDIR</td><td>A component of the path prefix is not a directory.</td></tr></table> <p>The <code>ftok()</code> function may fail if:</p> <table><tr><td>ENAMETOOLONG</td><td>Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>{PATH_MAX}</code>.</td></tr></table>	EACCES	Search permission is denied for a component of the path prefix.	ELOOP	Too many symbolic links were encountered in resolving <i>path</i> .	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> .	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.	ENAMETOOLONG	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>{PATH_MAX}</code> .
EACCES	Search permission is denied for a component of the path prefix.												
ELOOP	Too many symbolic links were encountered in resolving <i>path</i> .												
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> .												
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.												
ENAMETOOLONG	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>{PATH_MAX}</code> .												
USAGE	For maximum portability, <i>id</i> should be a single-byte character.												

Another way to compose keys is to include the project ID in the most significant byte and to use the remaining portion as a sequence number. There are many other ways to form keys, but it is necessary for each system to define standards for forming them. If some standard is not adhered to, it will be possible for unrelated processes to unintentionally interfere with each other's operation. It is still possible to interfere intentionally. Therefore, it is strongly suggested that the most significant byte of a key in some sense refer to a project so that keys do not conflict across a given system.

NOTES Since the `ftok()` function returns a value based on the *id* given and the file serial number of the file named by *path* in a type that is no longer large enough to hold all file serial numbers, it may return the same key for paths naming different files on large filesystems.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `msgget(2)`, `semget(2)`, `shmget(2)`, `stat(2)`, `attributes(5)`, `standards(5)`

ftw(3C)

NAME	ftw, nftw – walk a file tree																
SYNOPSIS	<pre>#include <ftw.h> int ftw(const char *path, int (*fn) (const char *, const struct stat *, int), int depth); int nftw(const char *path, int (*fn) (const char *, const struct stat *, int, struct FTW *), int depth, int flags);</pre>																
DESCRIPTION	<p>The <code>ftw()</code> function recursively descends the directory hierarchy rooted in <code>path</code>. For each object in the hierarchy, <code>ftw()</code> calls the user-defined function <code>fn</code>, passing it a pointer to a null-terminated character string containing the name of the object, a pointer to a <code>stat</code> structure (see <code>stat(2)</code>) containing information about the object, and an integer. Possible values of the integer, defined in the <code><ftw.h></code> header, are:</p> <table border="0"> <tr> <td style="padding-right: 20px;">FTW_F</td> <td>The object is a file.</td> </tr> <tr> <td>FTW_D</td> <td>The object is a directory.</td> </tr> <tr> <td>FTW_DNR</td> <td>The object is a directory that cannot be read. Descendants of the directory are not processed.</td> </tr> <tr> <td>FTW_NS</td> <td>The <code>stat()</code> function failed on the object because of lack of appropriate permission or the object is a symbolic link that points to a non-existent file. The <code>stat</code> buffer passed to <code>fn</code> is undefined.</td> </tr> </table> <p>The <code>ftw()</code> function visits a directory before visiting any of its descendants.</p> <p>The tree traversal continues until the tree is exhausted, an invocation of <code>fn</code> returns a non-zero value, or some error is detected within <code>ftw()</code> (such as an I/O error). If the tree is exhausted, <code>ftw()</code> returns 0. If <code>fn</code> returns a non-zero value, <code>ftw()</code> stops its tree traversal and returns whatever value was returned by <code>fn</code>.</p> <p>The <code>nftw()</code> function is similar to <code>ftw()</code> except that it takes the additional argument <code>flags</code>, which is a bitwise-inclusive OR of zero or more of the following flags:</p> <table border="0"> <tr> <td style="padding-right: 20px;">FTW_CHDIR</td> <td>If set, <code>nftw()</code> changes the current working directory to each directory as it reports files in that directory. If clear, <code>nftw()</code> does not change the current working directory.</td> </tr> <tr> <td>FTW_DEPTH</td> <td>If set, <code>nftw()</code> reports all files in a directory before reporting the directory itself. If clear, <code>nftw()</code> reports any directory before reporting the files in that directory.</td> </tr> <tr> <td>FTW_MOUNT</td> <td>If set, <code>nftw()</code> reports only files in the same file system as <code>path</code>. If clear, <code>nftw()</code> reports all files encountered during the walk.</td> </tr> <tr> <td>FTW_PHYS</td> <td>If set, <code>nftw()</code> performs a physical walk and does not follow symbolic links.</td> </tr> </table>	FTW_F	The object is a file.	FTW_D	The object is a directory.	FTW_DNR	The object is a directory that cannot be read. Descendants of the directory are not processed.	FTW_NS	The <code>stat()</code> function failed on the object because of lack of appropriate permission or the object is a symbolic link that points to a non-existent file. The <code>stat</code> buffer passed to <code>fn</code> is undefined.	FTW_CHDIR	If set, <code>nftw()</code> changes the current working directory to each directory as it reports files in that directory. If clear, <code>nftw()</code> does not change the current working directory.	FTW_DEPTH	If set, <code>nftw()</code> reports all files in a directory before reporting the directory itself. If clear, <code>nftw()</code> reports any directory before reporting the files in that directory.	FTW_MOUNT	If set, <code>nftw()</code> reports only files in the same file system as <code>path</code> . If clear, <code>nftw()</code> reports all files encountered during the walk.	FTW_PHYS	If set, <code>nftw()</code> performs a physical walk and does not follow symbolic links.
FTW_F	The object is a file.																
FTW_D	The object is a directory.																
FTW_DNR	The object is a directory that cannot be read. Descendants of the directory are not processed.																
FTW_NS	The <code>stat()</code> function failed on the object because of lack of appropriate permission or the object is a symbolic link that points to a non-existent file. The <code>stat</code> buffer passed to <code>fn</code> is undefined.																
FTW_CHDIR	If set, <code>nftw()</code> changes the current working directory to each directory as it reports files in that directory. If clear, <code>nftw()</code> does not change the current working directory.																
FTW_DEPTH	If set, <code>nftw()</code> reports all files in a directory before reporting the directory itself. If clear, <code>nftw()</code> reports any directory before reporting the files in that directory.																
FTW_MOUNT	If set, <code>nftw()</code> reports only files in the same file system as <code>path</code> . If clear, <code>nftw()</code> reports all files encountered during the walk.																
FTW_PHYS	If set, <code>nftw()</code> performs a physical walk and does not follow symbolic links.																

If `FTW_PHYS` is clear and `FTW_DEPTH` is set, `nftw()` follows links instead of reporting them, but does not report any directory that would be a descendant of itself. If `FTW_PHYS` is clear and `FTW_DEPTH` is clear, `nftw()` follows links instead of reporting them, but does not report the contents of any directory that would be a descendant of itself.

At each file it encounters, `nftw()` calls the user-supplied function `fn` with four arguments:

- The first argument is the pathname of the object.
- The second argument is a pointer to the `stat` buffer containing information on the object.
- The third argument is an integer giving additional information. Its value is one of the following:

<code>FTW_F</code>	The object is a file.
<code>FTW_D</code>	The object is a directory.
<code>FTW_DP</code>	The object is a directory and subdirectories have been visited. (This condition only occurs if the <code>FTW_DEPTH</code> flag is included in flags.)
<code>FTW_SL</code>	The object is a symbolic link. (This condition only occurs if the <code>FTW_PHYS</code> flag is included in flags.)
<code>FTW_SLN</code>	The object is a symbolic link that points to a non-existent file. (This condition only occurs if the <code>FTW_PHYS</code> flag is not included in flags.)
<code>FTW_DNR</code>	The object is a directory that cannot be read. The user-defined function <code>fn</code> will not be called for any of its descendants.
<code>FTW_NS</code>	The <code>stat()</code> function failed on the object because of lack of appropriate permission. The <code>stat</code> buffer passed to <code>fn</code> is undefined. Failure of <code>stat()</code> for any other reason is considered an error and <code>nftw()</code> returns <code>-1</code> .

- The fourth argument is a pointer to an `FTW` structure that contains the following members:

```
int    base;
int    level;
```

The `base` member is the offset of the object's filename in the pathname passed as the first argument to `fn()`. The value of `level` indicates the depth relative to the root of the walk, where the root level is 0.

The results are unspecified if the application-supplied `fn()` function does not preserve the current working directory.

ftw(3C)

Both `ftw()` and `nftw()` use one file descriptor for each level in the tree. The *depth* argument limits the number of file descriptors used. If *depth* is zero or negative, the effect is the same as if it were 1. It must not be greater than the number of file descriptors currently available for use. The `ftw()` function runs faster if *depth* is at least as large as the number of levels in the tree. When `ftw()` and `nftw()` return, they close any file descriptors they have opened; they do not close any file descriptors that might have been opened by *fn*.

RETURN VALUES

If the tree is exhausted, `ftw()` and `nftw()` return 0. If the function pointed to by *fn* returns a non-zero value, `ftw()` and `nftw()` stop their tree traversal and return whatever value was returned by the function pointed to by *fn*. If `ftw()` and `nftw()` detect an error, they return `-1` and set `errno` to indicate the error.

If `ftw()` and `nftw()` encounter an error other than `EACCES` (see `FTW_DNR` and `FTW_NS` above), they return `-1` and set `errno` to indicate the error. The external variable `errno` can contain any error value that is possible when a directory is opened or when one of the `stat` functions is executed on a directory or file.

ERRORS

The `ftw()` and `nftw()` functions will fail if:

<code>ELOOP</code>	A loop exists in symbolic links encountered during resolution of the <i>path</i> argument
<code>ENAMETOOLONG</code>	The length of the path exceeds <code>{PATH_MAX}</code> , or a path name component is longer than <code>{NAME_MAX}</code> .
<code>ENOENT</code>	A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.
<code>ENOTDIR</code>	A component of <i>path</i> is not a directory.
<code>EOVERFLOW</code>	A field in the <code>stat</code> structure cannot be represented correctly in the current programming environment for one or more files found in the file hierarchy.

The `ftw()` function will fail if:

<code>EACCES</code>	Search permission is denied for any component of <i>path</i> or read permission is denied for <i>path</i> .
---------------------	---

The `nftw()` function will fail if:

<code>EACCES</code>	Search permission is denied for any component of <i>path</i> or read permission is denied for <i>path</i> , or <i>fn()</i> returns <code>-1</code> and does not reset <code>errno</code> .
---------------------	--

The `nftw()` and `ftw()` functions may fail if:

<code>ELOOP</code>	Too many symbolic links were encountered during resolution of the <i>path</i> argument.
<code>ENAMETOOLONG</code>	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>{PATH_MAX}</code> .

The `ftw()` function may fail if:

`EINVAL` The value of the *depth* argument is invalid.

The `nftw()` function may fail if:

`EMFILE` There are `{OPEN_MAX}` file descriptors currently open in the calling process.

`ENFILE` Too many files are currently open in the system.

If the function pointed to by *fn* encounters system errors, `errno` may be set accordingly.

EXAMPLES **EXAMPLE 1** Walk a directory structure using `ftw()`.

The following example walks the current directory structure, calling the *fn()* function for every directory entry, using at most 10 file descriptors:

```
#include <ftw.h>
...
if (ftw(".", fn, 10) != 0) {
    perror("ftw"); exit(2);
}
```

EXAMPLE 2 Walk a directory structure using `nftw()`.

The following example walks the `/tmp` directory and its subdirectories, calling the `nftw()` function for every directory entry, to a maximum of 5 levels deep.

```
#include <ftw.h>
...
int nftwfunc(const char *, const struct stat *, int, struct FTW *);
int nftwfunc(const char *filename, const struct stat *statptr,
             int fileflags, struct FTW *pftw)
{
    return 0;
}
...
char *startpath = "/tmp";
int depth = 5;
int flags = FTW_CHDIR | FTW_DEPTH | FTW_MOUNT;
int ret;
ret = nftw(startpath, nftwfunc, depth, flags);
```

USAGE Because `ftw()` and `nftw()` are recursive, they can terminate with a memory fault when applied by a thread with a small stack to very deep file structures.

ftw(3C)

The `ftw()` and `nftw()` functions allocate resources (memory, file descriptors) during their operation. If `ftw()` they are forcibly terminated, such as by `longjmp(3C)` being executed by `fn` or an interrupt routine, they will not have a chance to free those resources, so they remain permanently allocated. A safe way to handle interrupts is to store the fact that an interrupt has occurred and arrange to have `fn` return a non-zero value at its next invocation.

The `ftw()` and `nftw()` functions have transitional interfaces for 64-bit file offsets. See `lf64(5)`.

The `ftw()` function is safe in multithreaded applications. The `nftw()` function is safe in multithreaded applications when the `FTW_CHDIR` flag is not set.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

SEE ALSO `stat(2)`, `longjmp(3C)`, `attributes(5)`, `lf64(5)`, `standards(5)`

NAME	fwide – set stream orientation						
SYNOPSIS	<pre>#include <stdio.h> #include <wchar.h> int fwide(FILE *<i>stream</i>, int <i>mode</i>);</pre>						
DESCRIPTION	<p>The <code>fwide()</code> function determines the orientation of the stream pointed to by <i>stream</i>. If <i>mode</i> is greater than 0, the function first attempts to make the stream wide-orientated. If <i>mode</i> is less than 0, the function first attempts to make the stream byte-orientated. Otherwise, <i>mode</i> is 0 and the function does not alter the orientation of the stream.</p> <p>If the orientation of the stream has already been determined, <code>fwide()</code> does not change it.</p> <p>Because no return value is reserved to indicate an error, an application wishing to check for error situations should set <code>errno</code> to 0, then call <code>fwide()</code>, then check <code>errno</code> and if it is non-zero, assume an error has occurred.</p>						
RETURN VALUES	The <code>fwide()</code> function returns a value greater than 0 if, after the call, the stream has wide-orientation, a value less than 0 if the stream has byte-orientation, or 0 if the stream has no orientation.						
ERRORS	<p>The <code>fwide()</code> function may fail if:</p> <p><code>EBADF</code> The <i>stream</i> argument is not a valid stream.</p>						
USAGE	A call to <code>fwide()</code> with <i>mode</i> set to 0 can be used to determine the current orientation of a stream.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>attributes(5)</code> , <code>standards(5)</code>						

fwprintf(3C)

NAME	fwprintf, wprintf, swprintf – print formatted wide-character output
SYNOPSIS	<pre>#include <stdio.h> #include <wchar.h> int fwprintf(FILE *restrict <i>stream</i>, const wchar_t *restrict <i>format</i>, ...); int wprintf(const wchar_t *restrict <i>format</i>, ...); int swprintf(wchar_t *restrict <i>s</i>, size_t <i>n</i>, const wchar_t *restrict <i>format</i>, ...);</pre>
DESCRIPTION	<p>The <code>fwprintf()</code> function places output on the named output <i>stream</i>. The <code>wprintf()</code> function places output on the standard output stream <code>stdout</code>. The <code>swprintf()</code> function places output followed by the null wide-character in consecutive wide-characters starting at <i>*s</i>; no more than <i>n</i> wide-characters are written, including a terminating null wide-character, which is always added (unless <i>n</i> is zero).</p> <p>Each of these functions converts, formats and prints its arguments under control of the <i>format</i> wide-character string. The <i>format</i> is composed of zero or more directives: <i>ordinary wide-characters</i>, which are simply copied to the output stream and <i>conversion specifications</i>, each of which results in the fetching of zero or more arguments. The results are undefined if there are insufficient arguments for the <i>format</i>. If the <i>format</i> is exhausted while arguments remain, the excess arguments are evaluated but are otherwise ignored.</p> <p>Conversions can be applied to the <i>n</i>th argument after the <i>format</i> in the argument list, rather than to the next unused argument. In this case, the conversion wide-character <code>%</code> (see below) is replaced by the sequence <code>%n\$</code>, where <i>n</i> is a decimal integer in the range <code>[1, NL_ARGMAX]</code>, giving the position of the argument in the argument list. This feature provides for the definition of format wide-character strings that select arguments in an order appropriate to specific languages (see the <code>EXAMPLES</code> section).</p> <p>In format wide-character strings containing the <code>%n\$</code> form of conversion specifications, numbered arguments in the argument list can be referenced from the format wide-character string as many times as required.</p> <p>In format wide-character strings containing the <code>%</code> form of conversion specifications, each argument in the argument list is used exactly once.</p> <p>All forms of the <code>fwprintf()</code> functions allow for the insertion of a language-dependent radix character in the output string, output as a wide-character value. The radix character is defined in the program's locale (category <code>LC_NUMERIC</code>). In the POSIX locale, or in a locale where the radix character is not defined, the radix character defaults to a period (<code>.</code>).</p> <p>Each conversion specification is introduced by the <code>%</code> wide-character or by the wide-character sequence <code>%n\$</code>, after which the following appear in sequence:</p> <ul style="list-style-type: none">■ Zero or more <i>flags</i> (in any order), which modify the meaning of the conversion specification.

- An optional minimum *field width*. If the converted value has fewer wide-characters than the field width, it will be padded with spaces by default on the left; it will be padded on the right, if the left-adjustment flag (-), described below, is given to the field width. The field width takes the form of an asterisk (*), described below, or a decimal integer.
- An optional *precision* that gives the minimum number of digits to appear for the d, i, o, u, x, and X conversions; the number of digits to appear after the radix character for the a, A, e, E, f, and F conversions; the maximum number of significant digits for the g and G conversions; or the maximum number of wide-characters to be printed from a string in s conversions. The precision takes the form of a period (.) followed by either an asterisk (*), described below, or an optional decimal digit string, where a null digit string is treated as 0. If a precision appears with any other conversion wide-character, the behavior is undefined.
- An optional length modifier that specifies the size of the argument.
- A *conversion specifier* wide character that indicates the type of conversion to be applied.

A field width, or precision, or both, may be indicated by an asterisk (*). In this case an argument of type `int` supplies the field width or precision. Arguments specifying field width, or precision, or both must appear in that order before the argument, if any, to be converted. A negative field width is taken as a - flag followed by a positive field width. A negative precision is taken as if the precision were omitted. In format wide-character strings containing the `%n$` form of a conversion specification, a field width or precision may be indicated by the sequence `*m$`, where `m` is a decimal integer in the range `[1, NL_ARGMAX]` giving the position in the argument list (after the format argument) of an integer argument containing the field width or precision, for example:

```
wprintf(L"%1$d:%2$.*3$d:%4$.*3$d\n", hour, min, precision, sec);
```

The *format* can contain either numbered argument specifications (that is, `%n$` and `*m$`), or unnumbered argument specifications (that is, `%` and `*`), but normally not both. The only exception to this is that `%%` can be mixed with the `%n$` form. The results of mixing numbered and unnumbered argument specifications in a *format* wide-character string are undefined. When numbered argument specifications are used, specifying the *N*th argument requires that all the leading arguments, from the first to the (*N*-1)th, are specified in the format wide-character string.

The flag wide-characters and their meanings are:

- ' The integer portion of the result of a decimal conversion (`%i`, `%d`, `%u`, `%f`, `%F`, `%g`, or `%G`) will be formatted with thousands' grouping wide-characters. For other conversions the behavior is undefined. The non-monetary grouping wide-character is used.
- The result of the conversion will be left-justified within the field. The conversion will be right-justified if this flag is not specified.

fwprintf(3C)

+	The result of a signed conversion will always begin with a sign (+ or -). The conversion will begin with a sign only when a negative value is converted if this flag is not specified.
space	If the first wide-character of a signed conversion is not a sign or if a signed conversion results in no wide-characters, a space will be prefixed to the result. This means that if the space and + flags both appear, the space flag will be ignored.
#	This flag specifies that the value is to be converted to an alternative form. For o conversion, it increases the precision (if necessary) to force the first digit of the result to be 0. For x or X conversions, a non-zero result will have 0x (or 0X) prefixed to it. For a, A, e, E, f, F, g, or G conversions, the result will always contain a radix character, even if no digits follow it. Without this flag, a radix character appears in the result of these conversions only if a digit follows it. For g and G conversions, trailing zeros will <i>not</i> be removed from the result as they normally are. For other conversions, the behavior is undefined.
0	For d, i, o, u, x, X, a, A, e, E, f, F, g, and G conversions, leading zeros (following any indication of sign or base) are used to pad to the field width; no space padding is performed. If the 0 and - flags both appear, the 0 flag will be ignored. For d, i, o, u, x, and X conversions, if a precision is specified, the 0 flag will be ignored. If the 0 and ' flags both appear, the grouping wide-characters are inserted before zero padding. For other conversions, the behavior is undefined.
The length modifiers and their meanings:	
hh	Specifies that a following d, i, o, u, x, or X conversion specifier applies to a signed char or unsigned char argument (the argument will have been promoted according to the integer promotions, but its value shall be converted to signed char or unsigned char before printing); or that a following n conversion specifier applies to a pointer to a signed char argument.
h	Specifies that a following d, i, o, u, x, or X conversion specifier applies to a short or unsigned short argument (the argument will have been promoted according to the integer promotions, but its value shall be converted to short or unsigned short before printing); or that a following n conversion specifier applies to a pointer to a short argument.
l (ell)	Specifies that a following d, i, o, u, x, or X conversion specifier applies to a long or unsigned long argument; that a following n conversion specifier applies to a pointer to a long argument; that a following c conversion specifier applies to a wint_t

	argument; that a following <i>s</i> conversion specifier applies to a pointer to a <code>wchar_t</code> argument; or has no effect on a following <i>a</i> , <i>A</i> , <i>e</i> , <i>E</i> , <i>f</i> , <i>F</i> , <i>g</i> , or <i>G</i> conversion specifier.
ll (ell-ell)	Specifies that a following <i>d</i> , <i>i</i> , <i>o</i> , <i>u</i> , <i>x</i> , or <i>X</i> conversion specifier applies to a <code>long long</code> or unsigned <code>long long</code> argument; or that a following <i>n</i> conversion specifier applies to a pointer to a <code>long long</code> argument.
j	Specifies that a following <i>d</i> , <i>i</i> , <i>o</i> , <i>u</i> , <i>x</i> , or <i>X</i> conversion specifier applies to an <code>intmax_t</code> or <code>uintmax_t</code> argument; or that a following <i>n</i> conversion specifier applies to a pointer to an <code>intmax_t</code> argument.
z	Specifies that a following <i>d</i> , <i>i</i> , <i>o</i> , <i>u</i> , <i>x</i> , or <i>X</i> conversion specifier applies to a <code>size_t</code> or the corresponding signed integer type argument; or that a following <i>n</i> conversion specifier applies to a pointer to a signed integer type corresponding to <code>size_t</code> argument.
t	Specifies that a following <i>d</i> , <i>i</i> , <i>o</i> , <i>u</i> , <i>x</i> , or <i>X</i> conversion specifier applies to a <code>ptrdiff_t</code> or the corresponding unsigned type argument; or that a following <i>n</i> conversion specifier applies to a pointer to a <code>ptrdiff_t</code> argument.
L	Specifies that a following <i>a</i> , <i>A</i> , <i>e</i> , <i>E</i> , <i>f</i> , <i>F</i> , <i>g</i> , or <i>G</i> conversion specifier applies to a long double argument.

If a length modifier appears with any conversion specifier other than as specified above, the behavior is undefined.

The conversion wide-characters and their meanings are:

<i>d</i> , <i>i</i>	The <code>int</code> argument is converted to a signed decimal in the style <code>[-] dddd</code> . The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no wide-characters.
<i>o</i>	The unsigned <code>int</code> argument is converted to unsigned octal format in the style <code>dddd</code> . The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no wide-characters.
<i>u</i>	The unsigned <code>int</code> argument is converted to unsigned decimal format in the style <code>dddd</code> . The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no wide-characters.

fwprintf(3C)

x	<p>The unsigned <code>int</code> argument is converted to unsigned hexadecimal format in the style <i>ddd</i>; the letters <code>abcdef</code> are used. The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no wide-characters.</p>		
X	<p>Behaves the same as the <code>x</code> conversion wide-character except that letters "ABCDEF" are used instead of "abcdef".</p>		
f, F	<p>The <code>double</code> argument is converted to decimal notation in the style <i>[-]ddd.ddd</i>, where the number of digits after the radix character (see setlocale(3C)) is equal to the precision specification. If the precision is missing it is taken as 6; if the precision is explicitly 0 and the <code>#</code> flag is not specified, no radix character appears. If a radix character appears, at least 1 digit appears before it. The converted value is rounded to fit the specified output format according to the prevailing floating point rounding direction mode. If the conversion is not exact, an inexact exception is raised.</p> <p>For the <code>f</code> specifier, a <code>double</code> argument representing an infinity or NaN is converted in the style of the <code>e</code> conversion specifier, except that for an infinite argument, "infinity" or "Infinity" is printed when the precision is at least 8 and "inf" or "Inf" is printed otherwise.</p> <p>For the <code>F</code> specifier, a <code>double</code> argument representing an infinity or NaN is converted in the SUSv3 style of the <code>E</code> conversion specifier, except that for an infinite argument, "INFINITY" is printed when the precision is at least 8 and or "INF" is printed otherwise.</p>		
e, E	<p>The <code>double</code> argument is converted in the style <i>[-]d.ddde ± dd</i>, where there is one digit before the radix character (which is non-zero if the argument is non-zero) and the number of digits after it is equal to the precision; if the precision is missing, it is taken as 6; if the precision is 0 and no <code>#</code> flag is present, no radix character appears. The converted value is rounded to fit the specified output format according to the prevailing floating point rounding direction mode. If the conversion is not exact, an inexact exception is raised. The <code>E</code> conversion wide-character will produce a number with <code>E</code> instead of <code>e</code> introducing the exponent. The exponent always contains at least two digits. If the value is 0, the exponent is 0.</p> <p>Infinity and NaN values are handled in one of the following ways:</p> <table><tr><td>SUSv3</td><td><p>For the <code>e</code> specifier, a <code>double</code> argument representing an infinity is printed as "[-]infinity", when the precision for the conversion is at least 7 and as "[-]inf" otherwise. A <code>double</code> argument representing a NaN is printed as "[-]nan". For the <code>E</code> specifier, "INF",</p></td></tr></table>	SUSv3	<p>For the <code>e</code> specifier, a <code>double</code> argument representing an infinity is printed as "[-]infinity", when the precision for the conversion is at least 7 and as "[-]inf" otherwise. A <code>double</code> argument representing a NaN is printed as "[-]nan". For the <code>E</code> specifier, "INF",</p>
SUSv3	<p>For the <code>e</code> specifier, a <code>double</code> argument representing an infinity is printed as "[-]infinity", when the precision for the conversion is at least 7 and as "[-]inf" otherwise. A <code>double</code> argument representing a NaN is printed as "[-]nan". For the <code>E</code> specifier, "INF",</p>		

	<p>"INFINITY", and "NaN" are printed instead of "inf", "infinity", and "nan", respectively. Printing of the sign follows the rules described above.</p>
Default	<p>A double argument representing an infinity is printed as "[−]Infinity", when the precision for the conversion is at least 7 and as "[−]Inf" otherwise. A double argument representing a NaN is printed as "[−]NaN". Printing of the sign follows the rules described above.</p>
g, G	<p>The double argument is converted in the style <code>f</code> or <code>e</code> (or in the style <code>E</code> in the case of a <code>G</code> conversion wide-character), with the precision specifying the number of significant digits. If an explicit precision is 0, it is taken as 1. The style used depends on the value converted; style <code>e</code> (or <code>E</code>) will be used only if the exponent resulting from such a conversion is less than -4 or greater than or equal to the precision. Trailing zeros are removed from the fractional portion of the result; a radix character appears only if it is followed by a digit.</p> <p>A double argument representing an infinity or NaN is converted in the style of the <code>e</code> or <code>E</code> conversion specifier, except that for an infinite argument, "infinity", "INFINITY", or "Infinity" is printed when the precision is at least 8 and "inf", "INF", or "Inf" is printed otherwise.</p>
a, A	<p>A double argument representing a floating-point number is converted in the style "[−]0xh.hhhhp±d", where the single hexadecimal digit preceding the radix point is 0 if the value converted is zero and 1 otherwise and the number of hexadecimal digits after it are equal to the precision; if the precision is missing, the number of digits printed after the radix point is 13 for the conversion of a double value, 16 for the conversion of a long double value on x86, and 28 for the conversion of a long double value on SPARC; if the precision is zero and the '#' flag is not specified, no decimal-point wide character appears. The letters "abcdef" are used for a conversion and the letters "ABCDEF" for A conversion. The A conversion specifier produces a number with 'X' and 'P' instead of 'x' and 'p'. The exponent always contains at least one digit, and only as many more digits as necessary to represent the decimal exponent of 2. If the value is zero, the exponent is zero.</p> <p>The converted value is rounded to fit the specified output format according to the prevailing floating point rounding direction mode. If the conversion is not exact, an inexact exception is raised.</p> <p>A double argument representing an infinity or NaN is converted in the SUSv3 style of an <code>e</code> or <code>E</code> conversion specifier.</p>

fwprintf(3C)

- c** If no `l` (ell) qualifier is present, the `int` argument is converted to a wide-character as if by calling the `btowc(3C)` function and the resulting wide-character is written. Otherwise the `wint_t` argument is converted to `wchar_t`, and written.
- s** If no `l` (ell) qualifier is present, the argument must be a pointer to a character array containing a character sequence beginning in the initial shift state. Characters from the array are converted as if by repeated calls to the `mbrtowc(3C)` function, with the conversion state described by an `mbstate_t` object initialized to zero before the first character is converted, and written up to (but not including) the terminating null wide-character. If the precision is specified, no more than that many wide-characters are written. If the precision is not specified or is greater than the size of the array, the array must contain a null wide-character.
- If an `l` (ell) qualifier is present, the argument must be a pointer to an array of type `wchar_t`. Wide characters from the array are written up to (but not including) a terminating null wide-character. If no precision is specified or is greater than the size of the array, the array must contain a null wide-character. If a precision is specified, no more than that many wide-characters are written.
- p** The argument must be a pointer to `void`. The value of the pointer is converted to a sequence of printable wide-characters.
- n** The argument must be a pointer to an integer into which is written the number of wide-characters written to the output so far by this call to one of the `fwprintf()` functions. No argument is converted.
- C** Same as `lc`.
- S** Same as `ls`.
- %** Output a `%` wide-character; no argument is converted. The entire conversion specification must be `%%`.

If a conversion specification does not match one of the above forms, the behavior is undefined.

In no case does a non-existent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is simply expanded to contain the conversion result. Characters generated by `fwprintf()` and `wprintf()` are printed as if `fputwc(3C)` had been called.

The `st_ctime` and `st_mtime` fields of the file will be marked for update between the call to a successful execution of `fwprintf()` or `wprintf()` and the next successful completion of a call to `fflush(3C)` or `fclose(3C)` on the same stream or a call to `exit(3C)` or `abort(3C)`.

RETURN VALUES Upon successful completion, these functions return the number of wide-characters transmitted excluding the terminating null wide-character in the case of `swprintf()` or a negative value if an output error was encountered.

If *n* or more wide characters were requested to be written, `swprintf()` returns a negative value.

ERRORS For the conditions under which `fwprintf()` and `wprintf()` will fail and may fail, refer to [fputwc\(3C\)](#).

In addition, all forms of `fwprintf()` may fail if:

EILSEQ A wide-character code that does not correspond to a valid character has been detected.

EINVAL There are insufficient arguments.

In addition, `wprintf()` and `fwprintf()` may fail if:

ENOMEM Insufficient storage space is available.

EXAMPLES **EXAMPLE 1** Print Language-dependent Date and Time Format.

To print the language-independent date and time format, the following statement could be used:

```
wprintf(format, weekday, month, day, hour, min);
```

For American usage, *format* could be a pointer to the wide-character string:

```
L"%s, %s %d, %d:%.2d\n"
```

producing the message:

```
Sunday, July 3, 10:02
```

whereas for German usage, *format* could be a pointer to the wide-character string:

```
L"%1$s, %3$d. %2$s, %4$d:%5$.2d\n"
```

producing the message:

```
Sonntag, 3. Juli, 10:02
```

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

SEE ALSO [btowc\(3C\)](#), [fputwc\(3C\)](#), [fwscanf\(3C\)](#), [mbrtowc\(3C\)](#), [setlocale\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NOTES The `fwprintf()`, `wprintf()`, and `swprintf()` functions can be used safely in multithreaded applications, as long as [setlocale\(3C\)](#) is not being called to change the locale.

fwprintf(3C)

If the `j` length modifier is used, 32-bit applications that were compiled using `c89` on releases prior to Solaris 10 will experience undefined behavior.

NAME	fwrite – binary output						
SYNOPSIS	<pre>#include <stdio.h> size_t fwrite(const void *ptr, size_t size, size_t nitems, FILE *stream);</pre>						
DESCRIPTION	<p>The <code>fwrite()</code> function writes, from the array pointed to by <code>ptr</code>, up to <code>nitems</code> elements whose size is specified by <code>size</code>, to the stream pointed to by <code>stream</code>. For each object, <code>size</code> calls are made to the <code>fputc(3C)</code> function, taking the values (in order) from an array of unsigned char exactly overlaying the object. The file-position indicator for the stream (if defined) is advanced by the number of bytes successfully written. If an error occurs, the resulting value of the file-position indicator for the stream is unspecified.</p> <p>The <code>st_ctime</code> and <code>st_mtime</code> fields of the file will be marked for update between the successful execution of <code>fwrite()</code> and the next successful completion of a call to <code>fflush(3C)</code> or <code>fclose(3C)</code> on the same stream or a call to <code>exit(2)</code> or <code>abort(3C)</code>.</p>						
RETURN VALUES	<p>The <code>fwrite()</code> function returns the number of elements successfully written, which might be less than <code>nitems</code> if a write error is encountered. If <code>size</code> or <code>nitems</code> is 0, <code>fwrite()</code> returns 0 and the state of the stream remains unchanged. Otherwise, if a write error occurs, the error indicator for the stream is set and <code>errno</code> is set to indicate the error.</p>						
ERRORS	Refer to <code>fputc(3C)</code> .						
USAGE	Because of possible differences in element length and byte ordering, files written using <code>fwrite()</code> are application-dependent, and possibly cannot be read using <code>fread(3C)</code> by a different application or by the same application on a different processor.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>write(2)</code> , <code>fclose(3C)</code> , <code>ferror(3C)</code> , <code>fopen(3C)</code> , <code>fread(3C)</code> , <code>getc(3C)</code> , <code>gets(3C)</code> , <code>printf(3C)</code> , <code>putc(3C)</code> , <code>puts(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

fwscanf(3C)

NAME	fwscanf, wscanf, swscanf, vfwscanf, vwscanf, vswscanf – convert formatted wide-character input
SYNOPSIS	<pre>#include <stdio.h> #include <wchar.h> int fwscanf(FILE *restrict <i>stream</i>, const wchar_t *restrict <i>format</i>, ...); int wscanf(const wchar_t *restrict <i>format</i>, ...); int swscanf(const wchar_t *restrict <i>s</i>, const wchar_t *restrict <i>format</i>, ...); #include <stdarg.h> #include <stdio.h> #include <wchar.h> int vfwscanf(FILE *restrict <i>stream</i>, const wchar_t *restrict <i>format</i>, va_list <i>arg</i>); int vwscanf(const wchar_t *restrict <i>ws</i>, const wchar_t *restrict <i>format</i>, va_list <i>arg</i>); int vswscanf(const wchar_t *restrict <i>format</i>, va_list <i>arg</i>);</pre>
DESCRIPTION	<p>The <code>fwscanf()</code> function reads from the named input <i>stream</i>.</p> <p>The <code>wscanf()</code> function reads from the standard input stream <code>stdin</code>.</p> <p>The <code>swscanf()</code> function reads from the wide-character string <i>s</i>.</p> <p>The <code>vfwscanf()</code>, <code>vswscanf()</code>, and <code>vwscanf()</code> functions are equivalent to the <code>fwscanf()</code>, <code>swscanf()</code>, and <code>wscanf()</code> functions, respectively, except that instead of being called with a variable number of arguments, they are called with an argument list as defined by the <code><stdarg.h></code> header. These functions do not invoke the <code>va_end()</code> macro. Applications using these functions should call <code>va_end(ap)</code> afterwards to clean up.</p> <p>Each function reads wide-characters, interprets them according to a format, and stores the results in its arguments. Each expects, as arguments, a control wide-character string <i>format</i> described below, and a set of <i>pointer</i> arguments indicating where the converted input should be stored. The result is undefined if there are insufficient arguments for the format. If the format is exhausted while arguments remain, the excess arguments are evaluated but are otherwise ignored.</p> <p>Conversions can be applied to the <i>n</i>th argument after the <i>format</i> in the argument list, rather than to the next unused argument. In this case, the conversion wide-character <code>%</code> (see below) is replaced by the sequence <code>%n\$</code>, where <i>n</i> is a decimal integer in the range <code>[1, NL_ARGMAX]</code>. This feature provides for the definition of format wide-character strings that select arguments in an order appropriate to specific languages. In format wide-character strings containing the <code>%n\$</code> form of conversion specifications, it is unspecified whether numbered arguments in the argument list can be referenced from the format wide-character string more than once.</p>

The *format* can contain either form of a conversion specification, that is, % or %n\$, but the two forms cannot normally be mixed within a single *format* wide-character string. The only exception to this is that %% or %* can be mixed with the %n\$ form.

The `fwscanf()` function in all its forms allows for detection of a language-dependent radix character in the input string, encoded as a wide-character value. The radix character is defined in the program's locale (category `LC_NUMERIC`). In the POSIX locale, or in a locale where the radix character is not defined, the radix character defaults to a period (.).

The format is a wide-character string composed of zero or more directives. Each directive is composed of one of the following: one or more white-space wide-characters (space, tab, newline, vertical-tab or form-feed characters); an ordinary wide-character (neither % nor a white-space character); or a conversion specification. Each conversion specification is introduced by a % or the sequence %n\$ after which the following appear in sequence:

- An optional assignment-suppressing character *.
- An optional non-zero decimal integer that specifies the maximum field width.
- An option length modifier that specifies the size of the receiving object.
- A conversion specifier wide-character that specifies the type of conversion to be applied. The valid conversion wide-characters are described below.

The `fwscanf()` functions execute each directive of the format in turn. If a directive fails, as detailed below, the function returns. Failures are described as input failures (due to the unavailability of input bytes) or matching failures (due to inappropriate input).

A directive composed of one or more white-space wide-characters is executed by reading input until no more valid input can be read, or up to the first wide-character which is not a white-space wide-character, which remains unread.

A directive that is an ordinary wide-character is executed as follows. The next wide-character is read from the input and compared with the wide-character that comprises the directive; if the comparison shows that they are not equivalent, the directive fails, and the differing and subsequent wide-characters remain unread.

A directive that is a conversion specification defines a set of matching input sequences, as described below for each conversion wide-character. A conversion specification is executed in the following steps:

Input white-space wide-characters (as specified by `iswspace(3C)`) are skipped, unless the conversion specification includes a [, c, or n conversion character.

An item is read from the input unless the conversion specification includes an n conversion wide-character. The length of the item read is limited to any specified maximum field width. In Solaris default mode, the input item is defined as the longest sequence of input wide-characters that forms a matching sequence. In some cases, `fwscanf()` might need to read several extra wide-characters beyond the end of the

fwscanf(3C)

input item to find the end of a matching sequence. In C99/SUSv3 mode, the input item is defined as the longest sequence of input wide-characters that is, or is a prefix of, a matching sequence. With this definition, `fwscanf()` need only read at most one wide-character beyond the end of the input item. Therefore, in C99/SUSv3 mode, some sequences that are acceptable to `wcstod(3C)`, `wcstol(3C)`, and similar functions are unacceptable to `fwscanf()`. In either mode, `fwscanf()` attempts to push back any excess bytes read using `ungetc(3C)`. Assuming all such attempts succeed, the first wide-character, if any, after the input item remains unread. If the length of the input item is 0, the conversion fails. This condition is a matching failure unless end-of-file, an encoding error, or a read error prevented input from the stream, in which case it is an input failure.

Except in the case of a `%` conversion wide-character, the input item (or, in the case of a `%n` conversion specification, the count of input wide-characters) is converted to a type appropriate to the conversion wide-character. If the input item is not a matching sequence, the execution of the conversion specification fails; this condition is a matching failure. Unless assignment suppression was indicated by a `*`, the result of the conversion is placed in the object pointed to by the first argument following the *format* argument that has not already received a conversion result if the conversion specification is introduced by `%`, or in the *n*th argument if introduced by the wide-character sequence `%n$`. If this object does not have an appropriate type, or if the result of the conversion cannot be represented in the space provided, the behavior is undefined.

The length modifiers and their meanings are:

hh	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to signed char or unsigned char.
h	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to short or unsigned short.
l (ell)	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to long or unsigned long; that a following <code>a</code> , <code>A</code> , <code>e</code> , <code>E</code> , <code>f</code> , <code>F</code> , <code>g</code> , or <code>G</code> conversion specifier applies to an argument with type pointer to double; or that a following <code>c</code> , <code>s</code> , or <code>[]</code> conversion specifier applies to an argument with type pointer to <code>wchar_t</code> .
ll (ell-ell)	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to long long or unsigned long long.
j	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to <code>intmax_t</code> or <code>uintmax_t</code> .
z	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to <code>size_t</code> or the corresponding signed integer type.
t	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to <code>ptrdiff_t</code> or the corresponding unsigned type.

- L Specifies that a following *a*, *A*, *e*, *E*, *f*, *F*, *g*, or *G* conversion specifier applies to an argument with type pointer to long double.

If a length modifier appears with any conversion specifier other than as specified above, the behavior is undefined.

The following conversion wide-characters are valid:

- d Matches an optionally signed decimal integer, whose format is the same as expected for the subject sequence of `wcstol(3C)` with the value 10 for the *base* argument. In the absence of a size modifier, the corresponding argument must be a pointer to `int`.
- i Matches an optionally signed integer, whose format is the same as expected for the subject sequence of `wcstol(3C)` with 0 for the *base* argument. In the absence of a size modifier, the corresponding argument must be a pointer to `int`.
- o Matches an optionally signed octal integer, whose format is the same as expected for the subject sequence of `wcstoul(3C)` with the value 8 for the *base* argument. In the absence of a size modifier, the corresponding argument must be a pointer to unsigned `int`.
- u Matches an optionally signed decimal integer, whose format is the same as expected for the subject sequence of `wcstoul(3C)` with the value 10 for the *base* argument. In the absence of a size modifier, the corresponding argument must be a pointer to unsigned `int`.
- x Matches an optionally signed hexadecimal integer, whose format is the same as expected for the subject sequence of `wcstoul(3C)` with the value 16 for the *base* argument. In the absence of a size modifier, the corresponding argument must be a pointer to unsigned `int`.
- a,e,f,g Matches an optionally signed floating-point number, whose format is the same as expected for the subject sequence of `wcstod(3C)`. In the absence of a size modifier, the corresponding argument must be a pointer to `float`. The *e*, *f*, and *g* specifiers match hexadecimal floating point values only in C99/SUSv3 (see `standards(5)`) mode, but the *a* specifier always matches hexadecimal floating point values.

These conversion specifiers match any subject sequence accepted by `strtod(3C)`, including the INF, INFINITY, NAN, and NAN(*n-char-sequence*) forms. The result of the conversion is the same as that of calling `strtod()` (or `strtof()` or `strtold()`) with the matching sequence, including the raising of floating point exceptions and the setting of `errno` to ERANGE, if applicable.

- s Matches a sequence of non white-space wide-characters. If no `l` (`ell`) qualifier is present, characters from the input field are converted as if by repeated calls to the `wcrtomb(3C)` function, with the conversion state described by an `mbstate_t` object initialized to zero before the first

wide-character is converted. The corresponding argument must be a pointer to a character array large enough to accept the sequence and the terminating null character, which will be added automatically.

Otherwise, the corresponding argument must be a pointer to an array of `wchar_t` large enough to accept the sequence and the terminating null wide-character, which will be added automatically.

- [Matches a non-empty sequence of wide-characters from a set of expected wide-characters (the *scanset*). If no `l` (ell) qualifier is present, wide-characters from the input field are converted as if by repeated calls to the `wcrtomb()` function, with the conversion state described by an `mbstate_t` object initialized to zero before the first wide-character is converted. The corresponding argument must be a pointer to a character array large enough to accept the sequence and the terminating null character, which will be added automatically.

If an `l` (ell) qualifier is present, the corresponding argument must be a pointer to an array of `wchar_t` large enough to accept the sequence and the terminating null wide-character, which will be added automatically.

The conversion specification includes all subsequent `widw` characters in the *format* string up to and including the matching right square bracket (`]`). The wide-characters between the square brackets (the *scanlist*) comprise the *scanset*, unless the wide-character after the left square bracket is a circumflex (`^`), in which case the *scanset* contains all wide-characters that do not appear in the *scanlist* between the circumflex and the right square bracket. If the conversion specification begins with `[]` or `[^]`, the right square bracket is included in the *scanlist* and the next right square bracket is the matching right square bracket that ends the conversion specification; otherwise the first right square bracket is the one that ends the conversion specification. If a minus-sign (`-`) is in the *scanlist* and is not the first wide-character, nor the second where the first wide-character is a `^`, nor the last wide-character, it indicates a range of characters to be matched.

- c Matches a sequence of wide-characters of the number specified by the field width (1 if no field width is present in the conversion specification). If no `l` (ell) qualifier is present, wide-characters from the input field are converted as if by repeated calls to the `wcrtomb()` function, with the conversion state described by an `mbstate_t` object initialized to zero before the first wide-character is converted. The corresponding argument must be a pointer to a character array large enough to accept the sequence. No null character is added.

Otherwise, the corresponding argument must be a pointer to an array of `wchar_t` large enough to accept the sequence. No null wide-character is added.

p	Matches the set of sequences that is the same as the set of sequences that is produced by the %p conversion of the corresponding <code>fwprintf(3C)</code> functions. The corresponding argument must be a pointer to a pointer to <code>void</code> . If the input item is a value converted earlier during the same program execution, the pointer that results will compare equal to that value; otherwise the behavior of the %p conversion is undefined.
n	No input is consumed. The corresponding argument must be a pointer to the integer into which is to be written the number of wide-characters read from the input so far by this call to the <code>fwscanf()</code> functions. Execution of a %n conversion specification does not increment the assignment count returned at the completion of execution of the function.
C	Same as <code>lc</code> .
S	Same as <code>ls</code> .
%	Matches a single %; no conversion or assignment occurs. The complete conversion specification must be %%.

If a conversion specification is invalid, the behavior is undefined.

The conversion characters A, E, F, G, and X are also valid and behave the same as, respectively, a, e, f, g, and x.

If end-of-file is encountered during input, conversion is terminated. If end-of-file occurs before any wide-characters matching the current conversion specification (except for %n) have been read (other than leading white-space, where permitted), execution of the current conversion specification terminates with an input failure. Otherwise, unless execution of the current conversion specification is terminated with a matching failure, execution of the following conversion specification (if any) is terminated with an input failure.

Reaching the end of the string in `swscanf()` is equivalent to encountering end-of-file for `fwscanf()`.

If conversion terminates on a conflicting input, the offending input is left unread in the input. Any trailing white space (including newline) is left unread unless matched by a conversion specification. The success of literal matches and suppressed assignments is only directly determinable via the %n conversion specification.

The `fwscanf()` and `wscanf()` functions may mark the `st_atime` field of the file associated with *stream* for update. The `st_atime` field will be marked for update by the first successful execution of `fgetc(3C)`, `fgetwc(3C)`, `fgets(3C)`, `fgetws(3C)`, `fread(3C)`, `getc(3C)`, `getwc(3C)`, `getchar(3C)`, `getwchar(3C)`, `gets(3C)`, `fscanf(3C)` or `fwscanf()` using *stream* that returns data not supplied by a prior call to `ungetc(3C)`.

fwscanf(3C)

RETURN VALUES Upon successful completion, these functions return the number of successfully matched and assigned input items; this number can be 0 in the event of an early matching failure. If the input ends before the first matching failure or conversion, EOF is returned. If a read error occurs the error indicator for the stream is set, EOF is returned, and `errno` is set to indicate the error.

ERRORS For the conditions under which the `fwscanf()` functions will fail and may fail, refer to [fgetwc\(3C\)](#).

In addition, `fwscanf()` may fail if:

EILSEQ Input byte sequence does not form a valid character.

EINVAL There are insufficient arguments.

USAGE In format strings containing the `%` form of conversion specifications, each argument in the argument list is used exactly once.

EXAMPLES **EXAMPLE 1** `wscanf()` example

The call:

```
int i, n; float x; char name[50];
n = wscanf(L"%d%f%s", &i, &x, name);
```

with the input line:

```
25 54.32E-1 Hamster
```

will assign to `n` the value 3, to `i` the value 25, to `x` the value 5.432, and `name` will contain the string Hamster.

The call:

```
int i; float x; char name[50];
(void) wscanf(L"%2d%f*d %[0123456789]", &i, &x, name);
```

with input:

```
56789 0123 56a72
```

will assign 56 to `i`, 789.0 to `x`, skip 0123, and place the string 56\0 in `name`. The next call to [getchar\(3C\)](#) will return the character a.

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [fgetc\(3C\)](#), [fgets\(3C\)](#), [fgetwc\(3C\)](#), [fgetws\(3C\)](#), [fread\(3C\)](#), [fscanf\(3C\)](#), [fwprintf\(3C\)](#), [getc\(3C\)](#), [getchar\(3C\)](#), [gets\(3C\)](#), [getwc\(3C\)](#), [getwchar\(3C\)](#), [setlocale\(3C\)](#), [strtod\(3C\)](#), [wrtomb\(3C\)](#), [wcstod\(3C\)](#), [wcstol\(3C\)](#), [wcstoul\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	getcpuid, gethomelgroup – obtain information on scheduling decisions
SYNOPSIS	<pre>#include <sys/processor.h> processorid_t getcpuid(void); ushort_t gethomelgroup(void);</pre>
DESCRIPTION	<p>The <code>getcpuid()</code> function returns the processor ID on which the calling thread is currently executing.</p> <p>The <code>gethomelgroup()</code> function returns the home locality group ID of the calling thread.</p>
RETURN VALUES	See DESCRIPTION.
ERRORS	No errors are defined.
USAGE	Both the current CPU and the home locality group can change at any time.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	<code>getcpuid()</code> is Stable; <code>gethomelgroup()</code> is Obsolete.
MT-Level	MT-Safe

SEE ALSO	<code>psradm(1M)</code> , <code>psrinfo(1M)</code> , <code>psrset(1M)</code> , <code>p_online(2)</code> , <code>processor_bind(2)</code> , <code>processor_info(2)</code> , <code>pset_assign(2)</code> , <code>pset_bind(2)</code> , <code>pset_info(2)</code> , <code>meminfo(2)</code> , <code>lgrp_home(3LGRP)</code> , <code>sysconf(3C)</code> , <code>attributes(5)</code>
NOTES	The <code>gethomelgroup()</code> function is obsolete and might be removed in a future release. It has been replaced by <code>lgrp_home(3LGRP)</code> .

getcwd(3C)

NAME	getcwd – get pathname of current working directory								
SYNOPSIS	<pre>#include <unistd.h> char *getcwd(char *buf, size_t size);</pre>								
DESCRIPTION	<p>The <code>getcwd()</code> function places an absolute pathname of the current working directory in the array pointed to by <i>buf</i>, and returns <i>buf</i>. The pathname copied to the array contains no components that are symbolic links. The <i>size</i> argument is the size in bytes of the character array pointed to by <i>buf</i> and must be at least one greater than the length of the pathname to be returned.</p> <p>If <i>buf</i> is not a null pointer, the pathname is stored in the space pointed to by <i>buf</i>.</p> <p>If <i>buf</i> is a null pointer, <code>getcwd()</code> obtains <i>size</i> bytes of space using <code>malloc(3C)</code>. The pointer returned by <code>getcwd()</code> can be used as the argument in a subsequent call to <code>free()</code>.</p>								
RETURN VALUES	Upon successful completion, <code>getcwd()</code> returns the <i>buf</i> argument. Otherwise, the function returns a null pointer and sets <code>errno</code> to indicate the error.								
ERRORS	<p>The <code>getcwd()</code> function will fail if:</p> <table><tr><td><code>EINVAL</code></td><td>The <i>size</i> argument is equal to 0.</td></tr><tr><td><code>ERANGE</code></td><td>The <i>size</i> argument is greater than 0 and less than the length of the pathname plus 1.</td></tr></table> <p>The <code>getcwd()</code> function may fail if:</p> <table><tr><td><code>EACCES</code></td><td>A parent directory cannot be read to get its name.</td></tr><tr><td><code>ENOMEM</code></td><td>Insufficient storage space is available.</td></tr></table>	<code>EINVAL</code>	The <i>size</i> argument is equal to 0.	<code>ERANGE</code>	The <i>size</i> argument is greater than 0 and less than the length of the pathname plus 1.	<code>EACCES</code>	A parent directory cannot be read to get its name.	<code>ENOMEM</code>	Insufficient storage space is available.
<code>EINVAL</code>	The <i>size</i> argument is equal to 0.								
<code>ERANGE</code>	The <i>size</i> argument is greater than 0 and less than the length of the pathname plus 1.								
<code>EACCES</code>	A parent directory cannot be read to get its name.								
<code>ENOMEM</code>	Insufficient storage space is available.								
EXAMPLES	<p>EXAMPLE 1 Determine the absolute pathname of the current working directory.</p> <p>The following example returns a pointer to an array that holds the absolute pathname of the current working directory. The pointer is returned in the <i>ptr</i> variable, which points to the <i>buf</i> array where the pathname is stored.</p> <pre>#include <stdlib.h> #include <unistd.h> ... long size; char *buf; char *ptr; size = pathconf(".", _PC_PATH_MAX); if ((buf = (char *)malloc((size_t)size)) != NULL) ptr = getcwd(buf, (size_t)size); ...</pre>								

USAGE Applications should exercise care when using `chdir(2)` in conjunction with `getcwd()`. The current working directory is global to all threads within a process. If more than one thread calls `chdir()` to change the working directory, a subsequent call to `getcwd()` could produce unexpected results.

EXAMPLES **EXAMPLE 2** Printing the current working directory

The following example prints the current working directory.

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

main( )
{
    char *cwd;
    if ((cwd = getcwd(NULL, 64)) == NULL) {
        perror("pwd");
        exit(2);
    }
    (void)printf("%s\n", cwd);
    free(cwd); /* free memory allocated by getcwd() */
    return(0);
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `chdir(2)`, `malloc(3C)`, `attributes(5)`, `standards(5)`

getdate(3C)

NAME	getdate – convert user format date and time																														
SYNOPSIS	<pre>#include <time.h> struct tm *getdate(const char *string); extern int getdate_err;</pre>																														
DESCRIPTION	<p>The <code>getdate()</code> function converts user-definable date and/or time specifications pointed to by <i>string</i> to a <code>tm</code> structure. The <code>tm</code> structure is defined in the <code><time.h></code> header.</p> <p>User-supplied templates are used to parse and interpret the input string. The templates are text files created by the user and identified via the environment variable <code>DATMSK</code>. Each line in the template represents an acceptable date and/or time specification using conversion specifications similar to those used by <code>strptime(3C)</code> and <code>strptime(3C)</code>. Dates before 1902 and after 2037 are illegal. The first line in the template that matches the input specification is used for interpretation and conversion into the internal time format.</p>																														
Conversion Specifications	<p>The following conversion specifications are supported:</p> <table><tr><td><code>%%</code></td><td>Same as <code>%</code>.</td></tr><tr><td><code>%a</code></td><td>Locale's abbreviated weekday name.</td></tr><tr><td><code>%A</code></td><td>Locale's full weekday name.</td></tr><tr><td><code>%b</code></td><td>Locale's abbreviated month name.</td></tr><tr><td><code>%B</code></td><td>Locale's full month name.</td></tr><tr><td><code>%c</code></td><td>Locale's appropriate date and time representation.</td></tr><tr><td><code>%C</code></td><td>Century number (the year divided by 100 and truncated to an integer as a decimal number [1,99]); single digits are preceded by 0; see <code>standards(5)</code>. If used without the <code>%y</code> specifier, this format specifier will assume the current year offset in whichever century is specified. The only valid years are between 1902-2037.</td></tr><tr><td><code>%d</code></td><td>day of month [01,31]; leading zero is permitted but not required.</td></tr><tr><td><code>%D</code></td><td>Date as <code>%m/%d/%y</code>.</td></tr><tr><td><code>%e</code></td><td>Same as <code>%d</code>.</td></tr><tr><td><code>%h</code></td><td>Locale's abbreviated month name.</td></tr><tr><td><code>%H</code></td><td>Hour (24-hour clock) [0,23]; leading zero is permitted but not required.</td></tr><tr><td><code>%I</code></td><td>Hour (12-hour clock) [1,12]; leading zero is permitted but not required.</td></tr><tr><td><code>%j</code></td><td>Day number of the year [1,366]; leading zeros are permitted but not required.</td></tr><tr><td><code>%m</code></td><td>Month number [1,12]; leading zero is permitted but not required.</td></tr></table>	<code>%%</code>	Same as <code>%</code> .	<code>%a</code>	Locale's abbreviated weekday name.	<code>%A</code>	Locale's full weekday name.	<code>%b</code>	Locale's abbreviated month name.	<code>%B</code>	Locale's full month name.	<code>%c</code>	Locale's appropriate date and time representation.	<code>%C</code>	Century number (the year divided by 100 and truncated to an integer as a decimal number [1,99]); single digits are preceded by 0; see <code>standards(5)</code> . If used without the <code>%y</code> specifier, this format specifier will assume the current year offset in whichever century is specified. The only valid years are between 1902-2037.	<code>%d</code>	day of month [01,31]; leading zero is permitted but not required.	<code>%D</code>	Date as <code>%m/%d/%y</code> .	<code>%e</code>	Same as <code>%d</code> .	<code>%h</code>	Locale's abbreviated month name.	<code>%H</code>	Hour (24-hour clock) [0,23]; leading zero is permitted but not required.	<code>%I</code>	Hour (12-hour clock) [1,12]; leading zero is permitted but not required.	<code>%j</code>	Day number of the year [1,366]; leading zeros are permitted but not required.	<code>%m</code>	Month number [1,12]; leading zero is permitted but not required.
<code>%%</code>	Same as <code>%</code> .																														
<code>%a</code>	Locale's abbreviated weekday name.																														
<code>%A</code>	Locale's full weekday name.																														
<code>%b</code>	Locale's abbreviated month name.																														
<code>%B</code>	Locale's full month name.																														
<code>%c</code>	Locale's appropriate date and time representation.																														
<code>%C</code>	Century number (the year divided by 100 and truncated to an integer as a decimal number [1,99]); single digits are preceded by 0; see <code>standards(5)</code> . If used without the <code>%y</code> specifier, this format specifier will assume the current year offset in whichever century is specified. The only valid years are between 1902-2037.																														
<code>%d</code>	day of month [01,31]; leading zero is permitted but not required.																														
<code>%D</code>	Date as <code>%m/%d/%y</code> .																														
<code>%e</code>	Same as <code>%d</code> .																														
<code>%h</code>	Locale's abbreviated month name.																														
<code>%H</code>	Hour (24-hour clock) [0,23]; leading zero is permitted but not required.																														
<code>%I</code>	Hour (12-hour clock) [1,12]; leading zero is permitted but not required.																														
<code>%j</code>	Day number of the year [1,366]; leading zeros are permitted but not required.																														
<code>%m</code>	Month number [1,12]; leading zero is permitted but not required.																														

%M	Minute [0,59]; leading zero is permitted but not required.
%n	Any white space.
%p	Locale's equivalent of either a.m. or p.m.
%r	Appropriate time representation in the 12-hour clock format with %p.
%R	Time as %H:%M.

SUSv3

%S	Seconds [0,60]; leading zero is permitted but not required. The range of values is [00,60] rather than [00,59] to allow for the occasional leap second.
----	---

Default and other standards

%S	Seconds [0,61]; leading zero is permitted but not required. The range of values is [00,61] rather than [00,59] to allow for the occasional leap second and even more occasional double leap second.
%t	Any white space.
%T	Time as %H:%M:%S.
%U	Week number of the year as a decimal number [0,53], with Sunday as the first day of the week; leading zero is permitted but not required.
%w	Weekday as a decimal number [0,6], with 0 representing Sunday.
%W	Week number of the year as a decimal number [0,53], with Monday as the first day of the week; leading zero is permitted but not required.
%x	Locale's appropriate date representation.
%X	Locale's appropriate time representation.
%y	Year within century. When a century is not otherwise specified, values in the range 69-99 refer to years in the twentieth century (1969 to 1999 inclusive); values in the range 00-68 refer to years in the twenty-first century (2000 to 2068 inclusive).
%Y	Year, including the century (for example, 1993).
%Z	Time zone name or no characters if no time zone exists.

Modified Conversion Specifications

Some conversion specifications can be modified by the E and O modifier characters to indicate that an alternative format or specification should be used rather than the one normally used by the unmodified specification. If the alternative format or specification does not exist in the current locale, the behavior be as if the unmodified conversion specification were used.

%Ec	Locale's alternative appropriate date and time representation.
%EC	Name of the base year (period) in the locale's alternative representation.
%Ex	Locale's alternative date representation.

getdate(3C)

%EX	Locale's alternative time representation.
%Ey	Offset from %EC (year only) in the locale's alternative representation.
%EY	Full alternative year representation.
%Od	Day of the month using the locale's alternative numeric symbols; leading zeros are permitted but not required.
%Oe	Same as %Od.
%OH	Hour (24-hour clock) using the locale's alternative numeric symbols.
%OI	Hour (12-hour clock) using the locale's alternative numeric symbols.
%Om	Month using the locale's alternative numeric symbols.
%OM	Minutes using the locale's alternative numeric symbols.
%OS	Seconds using the locale's alternative numeric symbols.
%OU	Week number of the year (Sunday as the first day of the week) using the locale's alternative numeric symbols.
%Ow	Number of the weekday (Sunday=0) using the locale's alternative numeric symbols.
%OW	Week number of the year (Monday as the first day of the week) using the locale's alternative numeric symbols.
%Oy	Year (offset from %C) in the locale's alternative representation and using the locale's alternative numeric symbols.

Internal Format Conversion

The following rules are applied for converting the input specification into the internal format:

- If only the weekday is given, today is assumed if the given day is equal to the current day and next week if it is less.
- If only the month is given, the current month is assumed if the given month is equal to the current month and next year if it is less and no year is given. (The first day of month is assumed if no day is given.)
- If only the year is given, the values of the `tm_mon`, `tm_mday`, `tm_yday`, `tm_wday`, and `tm_isdst` members of the returned `tm` structure are not specified.
- If the century is given, but the year within the century is not given, the current year within the century is assumed.
- If no hour, minute, and second are given, the current hour, minute, and second are assumed.
- If no date is given, today is assumed if the given hour is greater than the current hour and tomorrow is assumed if it is less.

General Specifications

A conversion specification that is an ordinary character is executed by scanning the next character from the buffer. If the character scanned from the buffer differs from the one comprising the conversion specification, the specification fails, and the differing and subsequent characters remain unscanned.

A series of conversion specifications composed of %n, %t, white space characters, or any combination is executed by scanning up to the first character that is not white space (which remains unscanned), or until no more characters can be scanned.

Any other conversion specification is executed by scanning characters until a character matching the next conversion specification is scanned, or until no more characters can be scanned. These characters, except the one matching the next conversion specification, are then compared to the locale values associated with the conversion specifier. If a match is found, values for the appropriate *tm* structure members are set to values corresponding to the locale information. If no match is found, `getdate()` fails and no more characters are scanned.

The month names, weekday names, era names, and alternative numeric symbols can consist of any combination of upper and lower case letters. The user can request that the input date or time specification be in a specific language by setting the `LC_TIME` category using `setlocale(3C)`.

RETURN VALUES

If successful, `getdate()` returns a pointer to a *tm* structure; otherwise, it returns `NULL` and sets the global variable `getdate_err` to indicate the error. Subsequent calls to `getdate()` alter the contents of `getdate_err`.

The following is a complete list of the `getdate_err` settings and their meanings:

- 1 The `DATEMSK` environment variable is null or undefined.
- 2 The template file cannot be opened for reading.
- 3 Failed to get file status information.
- 4 The template file is not a regular file.
- 5 An error is encountered while reading the template file.
- 6 The `malloc()` function failed (not enough memory is available).
- 7 There is no line in the template that matches the input.
- 8 The input specification is invalid (for example, February 31).

USAGE

The `getdate()` function makes explicit use of macros described on the `ctype(3C)` manual page.

EXAMPLES

EXAMPLE 1 Examples of the `getdate()` function.

The following example shows the possible contents of a template:

getdate(3C)

EXAMPLE 1 Examples of the `getdate()` function. (Continued)

```
%m
%A %B %d %Y, %H:%M:%S
%A
%B
%m/%d/%y %I %p
%d, %m, %Y %H:%M
at %A the %dst of %B in %Y
run job at %I %p, %B %dnd
%A den %d. %B %Y %H.%M Uhr
```

The following are examples of valid input specifications for the above template:

```
getdate("10/1/87 4 PM")
getdate("Friday")
getdate("Friday September 19 1987, 10:30:30")
getdate("24,9,1986 10:30")
getdate("at monday the 1st of december in 1986")
getdate("run job at 3 PM, december 2nd")
```

If the `LANG` environment variable is set to `de` (German), the following is valid:

```
getdate("freitag den 10. oktober 1986 10.30 Uhr")
```

Local time and date specification are also supported. The following examples show how local date and time specification can be defined in the template.

Invocation	Line in Template
<code>getdate("11/27/86")</code>	<code>%m/%d/%y</code>
<code>getdate("27.11.86")</code>	<code>%d.%m.%y</code>
<code>getdate("86-11-27")</code>	<code>%y-%m-%d</code>
<code>getdate("Friday 12:00:00")</code>	<code>%A %H:%M:%S</code>

The following examples illustrate the Internal Format Conversion rules. Assume that the current date is `Mon Sep 22 12:19:47 EDT 1986` and the `LANG` environment variable is not set.

Input	Template Line	Date
<code>Mon</code>	<code>%a</code>	<code>Mon Sep 22 12:19:48 EDT 1986</code>
<code>Sun</code>	<code>%a</code>	<code>Sun Sep 28 12:19:49 EDT 1986</code>
<code>Fri</code>	<code>%a</code>	<code>Fri Sep 26 12:19:49 EDT 1986</code>
<code>September</code>	<code>%B</code>	<code>Mon Sep 1 12:19:49 EDT 1986</code>

January	%B	Thu Jan 1 12:19:49 EST 1987
December	%B	Mon Dec 1 12:19:49 EDT 1986
Sep Mon	%b %a	Mon Sep 1 12:19:50 EDT 1986
Jan Fri	%b %a	Fri Jan 2 12:19:50 EST 1987
Dec Mon	%b %a	Mon Dec 1 12:19:50 EST 1986
Jan Wed 1989	%b %a %Y	Wed Jan 4 12:19:51 EST 1989
Fri 9	%a %H	Fri Sep 26 09:00:00 EDT 1986
Feb 10:30	%b %H:%S	Sun Feb 1 10:00:30 EST 1987
10:30	%H:%M	Tue Sep 23 10:30:00 EDT 1986
13:30	%H:%M	Mon Sep 22 13:30:00 EDT 1986

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `ctype(3C)`, `mktime(3C)`, `setlocale(3C)`, `strftime(3C)`, `strptime(3C)`, `attributes(5)`, `environ(5)`, `standards(5)`

getdtablesize(3C)

NAME	getdtablesize – get the file descriptor table size
SYNOPSIS	<pre>#include <unistd.h> int getdtablesize(void);</pre>
DESCRIPTION	The <code>getdtablesize()</code> function is equivalent to <code>getrlimit(2)</code> with the <code>RLIMIT_NOFILE</code> option.
RETURN VALUES	The <code>getdtablesize()</code> function returns the current soft limit as if obtained from a call to <code>getrlimit()</code> with the <code>RLIMIT_NOFILE</code> option.
ERRORS	No errors are defined.
USAGE	<p>There is no direct relationship between the value returned by <code>getdtablesize()</code> and <code>OPEN_MAX</code> defined in <code><limits.h></code>.</p> <p>Each process has a file descriptor table which is guaranteed to have at least 20 slots. The entries in the descriptor table are numbered with small integers starting at 0. The <code>getdtablesize()</code> function returns the current maximum size of this table by calling the <code>getrlimit()</code> function.</p>
SEE ALSO	<code>close(2)</code> , <code>getrlimit(2)</code> , <code>open(2)</code> , <code>setrlimit(2)</code> , <code>select(3C)</code>

NAME	getenv – return value for environment name
SYNOPSIS	<pre>#include <stdlib.h> char *getenv(const char *name) ;</pre>
DESCRIPTION	The <code>getenv()</code> function searches the environment list (see <code>environ(5)</code>) for a string of the form <code>name=value</code> and, if the string is present, returns a pointer to the <code>value</code> in the current environment.
RETURN VALUES	If successful, <code>getenv()</code> returns a pointer to the <code>value</code> in the current environment; otherwise, it returns a null pointer.
USAGE	The <code>getenv()</code> function can be safely called from a multithreaded application. Care must be exercised when using both <code>getenv()</code> and <code>putenv(3C)</code> in a multithreaded application. These functions examine and modify the environment list, which is shared by all threads in an application. The system prevents the list from being accessed simultaneously by two different threads. It does not, however, prevent two threads from successively accessing the environment list using <code>getenv()</code> or <code>putenv(3C)</code> .
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Safe

SEE ALSO `exec(2)`, `putenv(3C)`, `attributes(5)`, `environ(5)`, `standards(5)`

getexecname(3C)

NAME	getexecname – return pathname of executable				
SYNOPSIS	<pre>#include <stdlib.h> const char *getexecname(void);</pre>				
DESCRIPTION	<p>The <code>getexecname()</code> function returns the pathname (the first argument of one of the <code>exec</code> family of functions; see <code>exec(2)</code>) of the executable that started the process.</p> <p>Normally this is an absolute pathname, as the majority of commands are executed by the shells that append the command name to the user's <code>PATH</code> components. If this is not an absolute path, the output of <code>getcwd(3C)</code> can be prepended to it to create an absolute path, unless the process or one of its ancestors has changed its root directory or current working directory since the last successful call to one of the <code>exec</code> family of functions.</p>				
RETURN VALUES	If successful, <code>getexecname()</code> returns a pointer to the executables pathname; otherwise, it returns 0.				
USAGE	<p>The <code>getexecname()</code> function obtains the executable pathname from the <code>AT_SUN_EXECNAME</code> aux vector. These vectors are made available to dynamically linked processes only.</p> <p>A successful call to one of the <code>exec</code> family of functions will always have <code>AT_SUN_EXECNAME</code> in the aux vector. The associated pathname is guaranteed to be less than or equal to <code>PATH_MAX</code>, not counting the trailing null byte that is always present.</p>				
ATTRIBUTES	See <code>attributes(5)</code> 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>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe				
SEE ALSO	<code>exec(2)</code> , <code>getcwd(3C)</code> , <code>attributes(5)</code>				

NAME	getgrnam, getgrnam_r, getgrent, getgrent_r, getgrgid, getgrgid_r, setgrent, endgrent, fgetgrent, fgetgrent_r – group database entry functions
SYNOPSIS	<pre>#include <grp.h> struct group *getgrnam(const char *name); struct group *getgrnam_r(const char *name, struct group *grp, char *buffer, int bufsize); struct group *getgrent(void); struct group *getgrent_r(struct group *grp, char *buffer, int bufsize); struct group *getgrgid(gid_t gid); struct group *getgrgid_r(gid_t gid, struct group *grp, char *buffer, int bufsize); void setgrent(void); void endgrent(void); struct group *fgetgrent(FILE *f); struct group *fgetgrent_r(FILE *f, struct group *grp, char *buffer, int bufsize);</pre>
Standard conforming	<pre>cc [flag...] file... -D_POSIX_PTHREAD_SEMANTICS [library...] int getgrnam_r(const char *name, struct group *grp, char *buffer, size_t bufsize, struct group **result); int getgrgid_r(gid_t gid, struct group *grp, char *buffer, size_t bufsize, struct group **result);</pre>
DESCRIPTION	<p>These functions are used to obtain entries describing user groups. Entries can come from any of the sources for group specified in the <code>/etc/nsswitch.conf</code> file (see <code>nsswitch.conf(4)</code>).</p> <p>The <code>getgrnam()</code> function searches the group database for an entry with the group name specified by the character string parameter <code>name</code>.</p> <p>The <code>getgrgid()</code> function searches the group database for an entry with the (numeric) group id specified by <code>gid</code>.</p> <p>The <code>setgrent()</code>, <code>getgrent()</code>, and <code>endgrent()</code> functions are used to enumerate group entries from the database.</p> <p>The <code>setgrent()</code> function effectively rewinds the group database to allow repeated searches. It sets (or resets) the enumeration to the beginning of the set of group entries. This function should be called before the first call to <code>getgrent()</code>.</p> <p>The <code>getgrent()</code> function returns a pointer to a structure containing the broken-out fields of an entry in the group database. When first called, <code>getgrent()</code> returns a pointer to a group structure containing the next group structure in the group database. Successive calls can be used to search the entire database.</p>

getgrnam(3C)

The `endgrent()` function can be called to close the group database and deallocate resources when processing is complete. It is permissible, though possibly less efficient, for the process to call more group functions after calling `endgrent()`.

The `fgetgrent()` function, unlike the other functions above, does not use `nsswitch.conf`. It reads and parses the next line from the stream *f*, which is assumed to have the format of the group file (see `group(4)`).

Reentrant Interfaces

The `getgrnam()`, `getgrgid()`, `getgrent()`, and `fgetgrent()` functions use thread-specific storage that is reused in each call to one of these functions by the same thread, making them safe to use but not recommended for multithreaded applications.

The parallel functions `getgrnam_r()`, `getgrgid_r()`, `getgrent_r()`, and `fgetgrent_r()` provide reentrant interfaces for these operations.

Each reentrant interface performs the same operation as its non-reentrant counterpart, named by removing the `_r` suffix. The reentrant interfaces, however, use buffers supplied by the caller to store returned results instead of using thread-specific data that can be overwritten by each call. They are safe for use in both single-threaded and multithreaded applications.

Each reentrant interface takes the same arguments as its non-reentrant counterpart, as well as the following additional parameters. The *grp* argument must be a pointer to a `struct group` structure allocated by the caller. On successful completion, the function returns the group entry in this structure. Storage referenced by the group structure is allocated from the memory provided with the *buffer* argument that is *bufsize* characters in size. The maximum size needed for this buffer can be determined with the `_SC_GETGR_R_SIZE_MAX` `sysconf(3C)` parameter. The standard-conforming versions place a pointer to the modified *grp* structure in the *result* parameter, instead of returning a pointer to this structure. A null pointer is returned at the location pointed to by *result* on error or if the requested entry is not found.

For enumeration in multithreaded applications, the position within the enumeration is a process-wide property shared by all threads. The `setgrent()` function can be used in a multithreaded application but resets the enumeration position for all threads. If multiple threads interleave calls to `getgrent_r()`, the threads will enumerate disjoint subsets of the group database. Like their non-reentrant counterparts, `getgrnam_r()` and `getgrgid_r()` leave the enumeration position in an indeterminate state.

group Structure

Group entries are represented by the `struct group` structure defined in `<grp.h>`:

```
struct group {
    char *gr_name;           /* the name of the group */
    char *gr_passwd;        /* the encrypted group password */
    gid_t gr_gid;          /* the numerical group ID */
    char **gr_mem;         /* vector of pointers to member names */
};
```

RETURN VALUES

The `getgrnam()`, `getgrnam_r()`, `getgrgid()`, and `getgrgid_r()` functions each return a pointer to a `struct group` if they successfully locate the requested entry. They return a null pointer if either the requested entry was not found or an error occurred. On error, `errno` is set to indicate the error. The standard-conforming functions `getgrnam_r()` and `getgrgid_r()` return 0 upon success or an error number in case of failure.

The `getgrent()`, `getgrent_r()`, `fgetgrent()`, and `fgetgrent_r()` functions each return a pointer to a `struct group` if they successfully enumerate an entry; otherwise they return a null pointer on end-of-file or error. On error, `errno` is set to indicate the error.

The `getgrnam()`, `getgrgid()`, `getgrent()`, and `fgetgrent()` functions use thread-specific data storage, so returned data must be copied before a subsequent call to any of these functions if the data are to be saved.

When the pointer returned by the reentrant functions `getgrnam_r()`, `getgrgid_r()`, `getgrent_r()`, and `fgetgrent_r()` is non-null, it is always equal to the `grp` pointer that was supplied by the caller.

Applications wishing to check for error situations should set `errno` to 0 before calling `getgrnam()`, `getgrnam_r()`, `getgrent()`, `getgrent_r()`, `getgrgid()`, `getgrgid_r()`, `fgetgrent()`, and `fgetgrent_r()`. If these functions return a null pointer and `errno` is non-zero, an error occurred.

ERRORS

The `getgrent_r()`, `fgetgrent()`, and `fgetgrent_r()` functions will fail if:

- | | |
|--------|---|
| EIO | An I/O error has occurred. |
| ERANGE | Insufficient storage was supplied by <i>buffer</i> and <i>bufsize</i> to contain the data to be referenced by the resulting <code>group</code> structure. |

The `getgrent_r()` function will fail if:

- | | |
|--------|--|
| EMFILE | There are {OPEN_MAX} file descriptors currently open in the calling process. |
| ENFILE | The maximum allowable number of files is currently open in the system. |

The `getgrnam()`, `getgrnam_r()`, `getgrgid()`, `getgrgid_r()`, and `getgrent()` functions may fail if:

- | | |
|--------|--|
| EINTR | A signal was caught during the operation. |
| EIO | An I/O error has occurred. |
| EMFILE | There are {OPEN_MAX} file descriptors currently open in the calling process. |
| ENFILE | The maximum allowable number of files is currently open in the system. |

getgrnam(3C)

The `getgrnam_r()` and `getgrgid_r()` functions may fail if:

ERANGE Insufficient storage was supplied by *buffer* and *bufsize* to contain the data to be referenced by the resulting group structure.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	<code>endgrent()</code> , <code>getgrent()</code> , <code>getgrgid()</code> , <code>getgrgid_r()</code> , <code>getgrnam()</code> , <code>getgrnam_r()</code> , and <code>setgrent()</code> are Standard.
MT-Level	See Reentrant Interfaces in DESCRIPTION.

SEE ALSO Intro(3), `getpwnam(3C)`, `group(4)`, `nsswitch.conf(4)`, `passwd(4)`, `attributes(5)`, `standards(5)`

NOTES When compiling multithreaded programs, see Intro(3).

Use of the enumeration interfaces `getgrent()` and `getgrent_r()` is discouraged; enumeration is supported for the group file, NIS, and NIS+, but in general is not efficient and might not be supported for all database sources. The semantics of enumeration are discussed further in `nsswitch.conf(4)`.

Previous releases allowed the use of "+" and "-" entries in `/etc/group` to selectively include and exclude entries from NIS. The primary usage of these entries is superseded by the name service switch, so the "+/-" form might not be supported in future releases.

If required, the "+/-" functionality can still be obtained for NIS by specifying `compat` as the source for `group`.

If the "+/-" functionality is required in conjunction with NIS+, specify both `compat` as the source for `group` and `nisplus` as the source for the pseudo-database `group_compat`. See `group(4)`, and `nsswitch.conf(4)` for details.

Solaris 2.4 and earlier releases provided definitions of the `getgrnam_r()` and `getgrgid_r()` functions as specified in POSIX.1c Draft 6. The final POSIX.1c standard changed the interface for these functions. Support for the Draft 6 interface is provided for compatibility only and might not be supported in future releases. New applications and libraries should use the standard-conforming interface.

For POSIX.1c-conforming applications, the `_POSIX_PTHREAD_SEMANTICS` and `_REENTRANT` flags are automatically turned on by defining the `_POSIX_C_SOURCE` flag with a value $\geq 199506L$.

NAME | gethostid – get an identifier for the current host

SYNOPSIS | `#include <unistd.h>`
| `long gethostid(void);`

DESCRIPTION | The `gethostid()` function returns the 32-bit identifier for the current host. This identifier is taken from the CPU board's ID PROM. It is not guaranteed to be unique.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO | `hostid(1)`, `sysinfo(2)`, `attributes(5)`, `standards(5)`

gethostname(3C)

NAME gethostname, sethostname – get or set name of current host

SYNOPSIS

```
#include <unistd.h>

int gethostname(char *name, int namelen);
int sethostname(char *name, int namelen);
```

DESCRIPTION

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

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

Host names are limited to `MAXHOSTNAMELEN` characters, currently 256, defined in the `<netdb.h>` header.

RETURN VALUES Upon successful completion, `gethostname()` and `sethostname()` return 0. Otherwise, they return `-1` and set `errno` to indicate the error.

ERRORS

The `gethostname()` and `sethostname()` functions will fail if:

- EFAULT** The `name` or `namelen` argument gave an invalid address.

The `sethostname()` function will fail if:

- EPERM** The `{PRIV_SYS_ADMIN}` privilege was not asserted in the effective set of the calling process.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `sysinfo(2)`, `uname(2)`, `gethostid(3C)`, `attributes(5)`, `standards(5)`

NAME	gethrtime, gethrvtime – get high resolution time				
SYNOPSIS	<pre>#include <sys/time.h> hrtime_t gethrtime(void); hrtime_t gethrvtime(void);</pre>				
DESCRIPTION	<p>The <code>gethrtime()</code> function returns the current high-resolution real time. Time is expressed as nanoseconds since some arbitrary time in the past; it is not correlated in any way to the time of day, and thus is not subject to resetting or drifting by way of <code>adjtime(2)</code> or <code>settimeofday(3C)</code>. The hi-res timer is ideally suited to performance measurement tasks, where cheap, accurate interval timing is required.</p> <p>The <code>gethrvtime()</code> function returns the current high-resolution LWP virtual time, expressed as total nanoseconds of execution time.</p> <p>The <code>gethrtime()</code> and <code>gethrvtime()</code> functions both return an <code>hrtime_t</code>, which is a 64-bit (long long) signed integer.</p>				
EXAMPLES	<p>The following code fragment measures the average cost of <code>getpid(2)</code>:</p> <pre>hrtime_t start, end; int i, iters = 100; start = gethrtime(); for (i = 0; i < iters; i++) getpid(); end = gethrtime(); printf("Avg getpid() time = %lld nsec\n", (end - start) / iters);</pre>				
ATTRIBUTES	<p>See <code>attributes(5)</code> 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;">MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>proc(1)</code> , <code>adjtime(2)</code> , <code>gettimeofday(3C)</code> , <code>settimeofday(3C)</code> , <code>attributes(5)</code>				
NOTES	<p>Although the units of hi-res time are always the same (nanoseconds), the actual resolution is hardware dependent. Hi-res time is guaranteed to be monotonic (it won't go backward, it won't periodically wrap) and linear (it won't occasionally speed up or slow down for adjustment, like the time of day can), but not necessarily unique: two sufficiently proximate calls may return the same value.</p>				

getloadavg(3C)

NAME	getloadavg – get system load averages				
SYNOPSIS	<pre>#include <sys/loadavg.h> int getloadavg(double <i>loadavg</i>[], int <i>nelem</i>);</pre>				
DESCRIPTION	The <code>getloadavg()</code> function returns the number of processes in the system run queue averaged over various periods of time. Up to <i>nelem</i> samples are retrieved and assigned to successive elements of <i>loadavg</i> []. The system imposes a maximum of 3 samples, representing averages over the last 1, 5, and 15 minutes, respectively. The <code>LOADAVG_1MIN</code> , <code>LOADAVG_5MIN</code> , and <code>LOADAVG_15MIN</code> indices, defined in <code><sys/loadavg.h></code> , can be used to extract the data from the appropriate element of the <i>loadavg</i> [] array.				
RETURN VALUES	Upon successful completion, the number of samples actually retrieved is returned. If the load average was unobtainable, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.				
ERRORS	The <code>getloadavg()</code> function will fail if: <code>EINVAL</code> The number of elements specified is less than 0.				
USAGE	If the caller is in a non-global zone and the pools facility is active, the behavior of <code>getloadavg()</code> is equivalent to that of <code>pset_getloadavg(3C)</code> called with <i>psetid</i> set to <code>PS_MYID</code> .				
ATTRIBUTES	See <code>attributes(5)</code> 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	<code>uptime(1)</code> , <code>w(1)</code> , <code>pooladm(1M)</code> , <code>kstat(3PERL)</code> , <code>pset_getloadavg(3C)</code> , <code>standards(5)</code>				

NAME	getlogin, getlogin_r – get login name
SYNOPSIS	<pre>#include <unistd.h> char *getlogin(void); char *getlogin_r(char *name, int namelen);</pre>
Standard conforming	<pre>cc [flag ...] file... -D_POSIX_PTHREAD_SEMANTICS [library ...] int getlogin_r(char *name, size_t namesize);</pre>
DESCRIPTION	<p>The <code>getlogin()</code> function returns a pointer to the login name as found in <code>/var/adm/utmpx</code>. It can be used in conjunction with <code>getpwnam(3C)</code> to locate the correct password file entry when the same user ID is shared by several login names.</p> <p>If <code>getlogin()</code> is called within a process that is not attached to a terminal, it returns a null pointer. The correct procedure for determining the login name is to call <code>cuserid(3C)</code>, or to call <code>getlogin()</code> and if it fails to call <code>getpwuid(3C)</code>.</p> <p>The <code>getlogin_r()</code> function has the same functionality as <code>getlogin()</code> except that the caller must supply a buffer <code>name</code> with length <code>namelen</code> to store the result. The <code>name</code> buffer must be at least <code>_POSIX_LOGIN_NAME_MAX</code> bytes in size (defined in <code><limits.h></code>). The POSIX version (see <code>standards(5)</code>) of <code>getlogin_r()</code> takes a <code>namesize</code> parameter of type <code>size_t</code>.</p>
RETURN VALUES	<p>Upon successful completion, <code>getlogin()</code> returns a pointer to the login name or a null pointer if the user's login name cannot be found. Otherwise it returns a null pointer and sets <code>errno</code> to indicate the error.</p> <p>The standard-conforming <code>getlogin_r()</code> returns 0 if successful, or the error number upon failure.</p>
ERRORS	<p>The <code>getlogin_r()</code> function will fail if:</p> <ul style="list-style-type: none"> ERANGE The size of the buffer is smaller than the result to be returned. EINVAL And entry for the current user was not found in the <code>/var/adm/utmpx</code> file. <p>The <code>getlogin()</code> and <code>getlogin_r()</code> functions may fail if:</p> <ul style="list-style-type: none"> EMFILE There are <code>{OPEN_MAX}</code> file descriptors currently open in the calling process. ENFILE The maximum allowable number of files is currently open in the system. ENXIO The calling process has no controlling terminal. <p>The <code>getlogin_r()</code> function may fail if:</p> <ul style="list-style-type: none"> ERANGE The size of the buffer is smaller than the result to be returned.

getlogin(3C)

USAGE The return value of `getlogin()` points to thread-specific data whose content is overwritten on each call by the same thread.

Three names associated with the current process can be determined:
`getpwuid(geteuid())` returns the name associated with the effective user ID of the process; `getlogin()` returns the name associated with the current login activity; and `getpwuid(getuid())` returns the name associated with the real user ID of the process.

FILES `/var/adm/utmpx` user access and administration information

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	See below.

SEE ALSO `geteuid(2)`, `getuid(2)`, `cuserid(3C)`, `getgrnam(3C)`, `getpwnam(3C)`, `getpwuid(3C)`, `utmpx(4)`, `attributes(5)`, `standards(5)`

NOTES When compiling multithreaded programs, see `Intro(3)`.

The `getlogin()` function is safe to use in multithreaded applications, but is discouraged. The `getlogin_r()` function should be used instead.

Solaris 2.4 and earlier releases provided a `getlogin_r()` 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 standard-conforming interface.

NAME	getmntent, getmntany, getextmntent, hasmntopt, putmntent, resetmnttab – get mounted device information
SYNOPSIS	<pre>#include <stdio.h> #include <sys/mnttab.h> int getmntent(FILE *fp, struct mnttab *mp); int getmntany(FILE *fp, struct mnttab *mp, struct mnttab *mpref); int getextmntent(FILE *fp, struct extmnttab *mp, int len); char *hasmntopt(struct mnttab *mnt, char *opt); int putmntent(FILE *iop, struct mnttab *mp); void resetmnttab(FILE *fp);</pre>
DESCRIPTION	
getmntent() and getmntany()	<p>The <code>getmntent()</code> and <code>getmntany()</code> functions each fill in the structure pointed to by <code>mp</code> with the broken-out fields of a line in the <code>mnttab</code> file. Each line read from the file contains a <code>mnttab</code> structure, which is defined in the <code><sys/mnttab.h></code> header. The structure contains the following members, which correspond to the broken-out fields from a line in <code>/etc/mnttab</code> (see <code>mnttab(4)</code>).</p> <pre>char *mnt_special; /* name of mounted resource */ char *mnt_mountp; /* mount point */ char *mnt_fstype; /* type of file system mounted */ char *mnt_mntopts; /* options for this mount */ char *mnt_time; /* time file system mounted */</pre> <p>Fields with no actual content in <code>/etc/mnttab</code> are represented in the file as <code>"-"</code>. To clearly distinguish empty fields, <code>getmntent()</code> set the corresponding field in <code>mp</code> to <code>NULL</code>.</p> <p>Each <code>getmntent()</code> call causes a new line to be read from the <code>mnttab</code> file. Successive calls can be used to search the entire list. The <code>getmntany()</code> function searches the file referenced by <code>fp</code> until a match is found between a line in the file and <code>mpref</code>. A match occurs if all non-null entries in <code>mpref</code> match the corresponding fields in the file. These functions do not open, close, or rewind the file.</p>
getextmntent()	<p>The <code>getextmntent()</code> function is an extended version of the <code>getmntent()</code> function that returns, in addition to the information that <code>getmntent()</code> returns, the major and minor number of the mounted resource to which the line in <code>mnttab</code> corresponds. The <code>getextmntent()</code> function also fills in the <code>extmntent</code> structure defined in the <code><sys/mnttab.h></code> header. For <code>getextmntent()</code> to function properly, it must be notified when the <code>mnttab</code> file has been reopened or rewound since a previous <code>getextmntent()</code> call. This notification is accomplished by calling <code>resetmnttab()</code>. Otherwise, it behaves exactly as <code>getmntent()</code> described above.</p> <p>The data pointed to by the <code>mnttab</code> structure members are stored in a static area and must be copied to be saved between successive calls.</p>

getmntent(3C)

hasmntopt() The `hasmntopt()` function scans the `mnt_mntopts` member of the `mnttab` structure `mnt` for a substring that matches `opt`. It returns the address of the substring if a match is found; otherwise it returns 0. Substrings are delimited by commas and the end of the `mnt_mntopts` string.

putmntent() The `putmntent()` function is obsolete and no longer has any effect. Entries appear in `mnttab` as a side effect of a `mount(2)` call. The function name is still defined for transition purposes.

resetmnttab() The `resetmnttab()` function notifies `getextmntent()` to reload from the kernel the device information that corresponds to the new snapshot of the `mnttab` information (see `mnttab(4)`). Subsequent `getextmntent()` calls then return correct `extmnttab` information. This function should be called whenever the `mnttab` file is either rewound or closed and reopened before any calls are made to `getextmntent()`.

RETURN VALUES

getmntent() and **getmntany()** If the next entry is successfully read by `getmntent()` or a match is found with `getmntany()`, 0 is returned. If an EOF is encountered on reading, these functions return -1. If an error is encountered, a value greater than 0 is returned. The following error values are defined in `<sys/mnttab.h>`:

`MNT_TOOLONG` A line in the file exceeded the internal buffer size of `MNT_LINE_MAX`.

`MNT_TOOMANY` A line in the file contains too many fields.

`MNT_TOOFEW` A line in the file contains too few fields.

hasmntopt() Upon successful completion, `hasmntopt()` returns the address of the substring if a match is found. Otherwise, it returns 0.

putmntent() The `putmntent()` is obsolete and always returns -1.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `mount(2)`, `mnttab(4)`, `attributes(5)`

NAME	getnetgrent, getnetgrent_r, setnetgrent, endnetgrent, innetgr – get network group entry
SYNOPSIS	<pre>#include <netdb.h> int getnetgrent(char **<i>machinep</i>, char **<i>userp</i>, char **<i>domainp</i>); int getnetgrent_r(char **<i>machinep</i>, char **<i>userp</i>, char **<i>domainp</i>, char *<i>buffer</i>, int<i>buflen</i>); int setnetgrent(const char *<i>netgroup</i>); int endnetgrent(void); int innetgr(const char *<i>netgroup</i>, const char *<i>machine</i>, const char *<i>user</i>, const char *<i>domain</i>);</pre>
DESCRIPTION	<p>These functions are used to test membership in and enumerate members of “netgroup” network groups defined in a system database. Netgroups are sets of (machine,user,domain) triples (see netgroup(4)).</p> <p>These functions consult the source specified for netgroup in the /etc/nsswitch.conf file (see nsswitch.conf(4)).</p> <p>The function innetgr() returns 1 if there is a netgroup <i>netgroup</i> that contains the specified <i>machine</i>, <i>user</i>, <i>domain</i> triple as a member; otherwise it returns 0. Any of the supplied pointers <i>machine</i>, <i>user</i>, and <i>domain</i> may be NULL, signifying a “wild card” that matches all values in that position of the triple.</p> <p>The innetgr() function is safe for use in single-threaded and multithreaded applications.</p> <p>The functions setnetgrent(), getnetgrent(), and endnetgrent() are used to enumerate the members of a given network group.</p> <p>The function setnetgrent() establishes the network group specified in the parameter <i>netgroup</i> as the current group whose members are to be enumerated.</p> <p>Successive calls to the function getnetgrent() will enumerate the members of the group established by calling setnetgrent(); each call returns 1 if it succeeds in obtaining another member of the network group, or 0 if there are no further members of the group.</p> <p>When calling either getnetgrent() or getnetgrent_r(), addresses of the three character pointers are used as arguments, for example:</p> <pre>char *<i>mp</i>, *<i>up</i>, *<i>dp</i>; getnetgrent(&<i>mp</i>, &<i>up</i>, &<i>dp</i>);</pre> <p>Upon successful return from getnetgrent(), the pointer <i>mp</i> points to a string containing the name of the machine part of the member triple, <i>up</i> points to a string containing the user name and <i>dp</i> points to a string containing the domain name. If the pointer returned for <i>mp</i>, <i>up</i>, or <i>dp</i> is NULL, it signifies that the element of the netgroup contains wild card specifier in that position of the triple.</p>

getnetgrent(3C)

The pointers returned by `getnetgrent()` point into a buffer allocated by `setnetgrent()` that is reused by each call. This space is released when an `endnetgrent()` call is made, and should not be released by the caller. This implementation is not safe for use in multi-threaded applications.

The function `getnetgrent_r()` is similar to `getnetgrent()` function, but it uses a buffer supplied by the caller for the space needed to store the results. The parameter *buffer* should be a pointer to a buffer allocated by the caller and the length of this buffer should be specified by the parameter *buflen*. The buffer must be large enough to hold the data associated with the triple. The `getnetgrent_r()` function is safe for use both in single-threaded and multi-threaded applications.

The function `endnetgrent()` frees the space allocated by the previous `setnetgrent()` call. The equivalent of an `endnetgrent()` implicitly performed whenever a `setnetgrent()` call is made to a new network group.

Note that while `setnetgrent()` and `endnetgrent()` are safe for use in multi-threaded applications, the effect of each is process-wide. Calling `setnetgrent()` resets the enumeration position for all threads. If multiple threads interleave calls to `getnetgrent_r()` each will enumerate a disjoint subset of the netgroup. Thus the effective use of these functions in multi-threaded applications may require coordination by the caller.

ERRORS The function `getnetgrent_r()` will return 0 and set `errno` to `ERANGE` if the length of the buffer supplied by caller is not large enough to store the result. See [Intro\(2\)](#) for the proper usage and interpretation of `errno` in multi-threaded applications.

The functions `setnetgrent()` and `endnetgrent()` return 0 upon success.

FILES `/etc/nsswitch.conf`

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	See DESCRIPTION section.

SEE ALSO [Intro\(2\)](#), [Intro\(3\)](#), [netgroup\(4\)](#), [nsswitch.conf\(4\)](#), [attributes\(5\)](#)

WARNINGS The function `getnetgrent_r()` is included in this release on an uncommitted basis only, and is subject to change or removal in future minor releases.

NOTES Only the Network Information Services, NIS and NIS+, are supported as sources for the `netgroup` database.

When compiling multi-threaded applications, see [Intro\(3\)](#), *Notes On Multithread Applications*, for information about the use of the `_REENTRANT` flag.

NAME	getopt – command option parsing
SYNOPSIS	
SVID3, XPG3	<pre>#include <stdio.h> int getopt(int <i>argc</i>, char * const <i>argv</i>[], const char *<i>optstring</i>); extern char *<i>optarg</i>; extern int <i>optind</i>, <i>opterr</i>, <i>optopt</i>;</pre>
POSIX.2, XPG4, SUS, SUSv2	<pre>#include <unistd.h> int getopt(int <i>argc</i>, char * const <i>argv</i>[], const char *<i>optstring</i>); extern char *<i>optarg</i>; extern int <i>optind</i>, <i>opterr</i>, <i>optopt</i>;</pre>
DESCRIPTION	<p>The <code>getopt()</code> function is a command line parser that can be used by applications that follow Basic Utility Syntax Guidelines 3, 4, 5, 6, 7, 9, and 10 which parallel those defined by application portability standards (see <code>intro(1)</code>). It can also be used by applications which additionally follow the Command Line Interface Paradigm (CLIP) syntax extension guidelines 15, 16, and 17. It partially enforces guideline 18 by requiring that every option has a short-name, but it allows multiple long-names to be associated with an option. The remaining guidelines are not addressed by <code>getopt()</code> and are the responsibility of the application.</p> <p>The <code>argc</code> and <code>argv</code> arguments are the argument count and argument array as passed to <code>main</code> (see <code>exec(2)</code>). The <code>optstring</code> argument specifies the acceptable options. For utilities wanting to conform to the Basic Utility Syntax Guidelines, <code>optstring</code> is a string of recognized option characters. All option characters allowed by Utility Syntax Guideline 3 are allowed in <code>optstring</code>. If a character is followed by a colon (:), the option is expected to have an option-argument, which can be separated from it by white space. Utilities wanting to conform to the extended CLIP guidelines can specify long-option equivalents to short options by following the short-option character (and optional colon) with a sequence of strings, each enclosed in parentheses, that specify the long-option aliases.</p> <p>The <code>getopt()</code> function returns the short-option character in <code>optstring</code> that corresponds to the next option found in <code>argv</code>.</p> <p>The <code>getopt()</code> function places in <code>optind</code> the <code>argv</code> index of the next argument to be processed. The <code>optind</code> variable is external and is initialized to 1 before the first call to <code>getopt()</code>. The <code>getopt()</code> function sets the variable <code>optarg</code> to point to the start of the option-argument as follows:</p> <ul style="list-style-type: none"> ■ If the option is a short option and that character is the last character in the argument, then <code>optarg</code> contains the next element of <code>argv</code>, and <code>optind</code> is incremented by 2. ■ If the option is a short option and that character is not the last character in the argument, then <code>optarg</code> points to the string following the option character in that argument, and <code>optind</code> is incremented by 1.

getopt(3C)

- If the option is a long option and the character equals is not found in the argument, then *optarg* contains the next element of *argv*, and *optind* is incremented by 2.
- If the option is a long option and the character equals is found in the argument, then *optarg* points to the string following the equals character in that argument and *optind* is incremented by 1.

In all cases, if the resulting value of *optind* is not less than *argc*, this indicates a missing option-argument and `getopt()` returns an error indication.

When all options have been processed (that is, up to the first operand), `getopt()` returns -1. The special option "--" (two hyphens) can be used to delimit the end of the options; when it is encountered, -1 is returned and "--" is skipped. This is useful in delimiting non-option arguments that begin with "-" (hyphen).

If `getopt()` encounters a short-option character or a long-option string not described in the *optstring* argument, it returns the question-mark (?) character. If it detects a missing option-argument, it also returns the question-mark (?) character, unless the first character of the *optstring* argument was a colon (:), in which case `getopt()` returns the colon (:) character. For short options, `getopt()` sets the variable *optopt* to the option character that caused the error. For long options, *optopt* is set to the hyphen (-) character and the failing long option can be identified through *argv[optind-1]*. If the application has not set the variable *opterr* to 0 and the first character of *optstring* is not a colon (:), `getopt()` also prints a diagnostic message to *stderr*.

RETURN VALUES

The `getopt()` function returns the short-option character associated with the option recognized.

A colon (:) is returned if `getopt()` detects a missing argument and the first character of *optstring* was a colon (:).

A question mark (?) is returned if `getopt()` encounters an option not specified in *optstring* or detects a missing argument and the first character of *optstring* was not a colon (:).

Otherwise, `getopt()` returns -1 when all command line options are parsed.

ERRORS

No errors are defined.

EXAMPLES

EXAMPLE 1 Parsing Command Line Options

The following code fragment shows how you might process the arguments for a utility that can take the mutually-exclusive options a and b and the options f and o, both of which require arguments:

```
#include <unistd.h>

int
main(int argc, char *argv[ ])
{
    int c;
    int bflg, aflg, errflg;
```

EXAMPLE 1 Parsing Command Line Options (Continued)

```

char *ifile;
char *ofile;
extern char *optarg;
extern int optind, optopt;
. . .
while ((c = getopt(argc, argv, ":abf:o:")) != -1) {
    switch(c) {
        case 'a':
            if (bflg)
                errflg++;
            else
                aflg++;
            break;
        case 'b':
            if (aflg)
                errflg++;
            else {
                bflg++;
                bproc();
            }
            break;
        case 'f':
            ifile = optarg;
            break;
        case 'o':
            ofile = optarg;
            break;
        case ':': /* -f or -o without operand */
            fprintf(stderr,
                "Option -%c requires an operand\n", optopt);
            errflg++;
            break;
        case '?':
            fprintf(stderr,
                "Unrecognized option: -%c\n", optopt);
            errflg++;
    }
}
if (errflg) {
    fprintf(stderr, "usage: . . . ");
    exit(2);
}
for ( ; optind < argc; optind++) {
    if (access(argv[optind], R_OK)) {
        . . .
    }
}

```

This code accepts any of the following as equivalent:

```

cmd -ao arg path path
cmd -a -o arg path path
cmd -o arg -a path path
cmd -a -o arg -- path path
cmd -a -oarg path path
cmd -aoarg path path

```

getopt(3C)

EXAMPLE 2 Check Options and Arguments.

The following example parses a set of command line options and prints messages to standard output for each option and argument that it encounters.

```
#include <unistd.h>
#include <stdio.h>
...
int c;
char *filename;
extern char *optarg;
extern int optind, optopt, opterr;
...
while ((c = getopt(argc, argv, ":abf:")) != -1) {
    switch(c) {
        case 'a':
            printf("a is set\n");
            break;
        case 'b':
            printf("b is set\n");
            break;
        case 'f':
            filename = optarg;
            printf("filename is %s\n", filename);
            break;
        case ':':
            printf("-%c without filename\n", optopt);
            break;
        case '?':
            printf("unknown arg %c\n", optopt);
            break;
    }
}
```

This example can be expanded to be CLIP-compliant by substituting the following string for the optstring argument:

```
:a(ascii)b(binary)f:(in-file)o:(out-file)V(version)?(help)
```

and by replacing the '?' case processing with:

```
case 'V':
    fprintf(stdout, "cmd 1.1\n");
    exit(0);
case '?':
    if (optopt == '?') {
        print_help();
        exit(0);
    }
    if (optopt == '-')
        fprintf(stderr,
            "unrecognized option: %s\n", argv[optind-1]);
    else
        fprintf(stderr,
            "unrecognized option: -%c\n", optopt);
    errflg++;
    break;
```

and by replacing the ':' case processing with:

EXAMPLE 2 Check Options and Arguments. (Continued)

```
case ':': /* -f or -o without operand */
    if (optopt == '-')
        fprintf(stderr,
            "Option %s requires an operand\n", argv[optind-1]);
    else
        fprintf(stderr,
            "Option -%c requires an operand\n", optopt);
    errflg++;
    break;
```

While not encouraged by the CLIP specification, multiple long-option aliases can also be assigned as shown in the following example:

```
:a (ascii)b (binary) : (in-file) (input) o: (outfile) (output) V (version) ? (help)
```

ENVIRONMENT VARIABLES

See [environ\(5\)](#) for descriptions of the following environment variables that affect the execution of `getopt()`: `LANG`, `LC_ALL`, and `LC_MESSAGES`.

`LC_CTYPE` Determine the locale for the interpretation of sequences of bytes as characters in *optstring*.

USAGE

The `getopt()` function does not fully check for mandatory arguments because there is no unambiguous algorithm to do so. Given an option string `a:b` and the input `-a -b`, `getopt()` assumes that `-b` is the mandatory argument to the `-a` option and not that `-a` is missing a mandatory argument. Indeed, the only time a missing option-argument can be reliably detected is when the option is the final option on the command line and is not followed by any command arguments.

It is a violation of the Basic Utility Command syntax standard (see [intro\(1\)](#)) for options with arguments to be grouped with other options, as in `cmd -abo filename`, where `a` and `b` are options, `o` is an option that requires an argument, and `filename` is the argument to `o`. Although this syntax is permitted in the current implementation, it should not be used because it may not be supported in future releases. The correct syntax to use is:

```
cmd- ab -o filename
```

ATTRIBUTES

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	See below.
MT-Level	Unsafe

The Basic Utility Command syntax is Standard. The CLIP extensions are Evolving.

SEE ALSO

[intro\(1\)](#), [getopt\(1\)](#), [getopts\(1\)](#), [getsubopt\(3C\)](#), [gettext\(3C\)](#), [setlocale\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#), [standards\(5\)](#)

getopt_long(3C)

NAME	getopt_long, getopt_long_only, getopt_clip – parse long command options
SYNOPSIS	<pre>#include <getopt.h> int getopt_long(int <i>argc</i>, char * const <i>argv</i>, const char <i>*shortopts</i>, const struct option <i>*longopts</i>, int <i>*indexptr</i>); int getopt_long_only(int <i>argc</i>, char * const <i>argv</i>, const char <i>*shortopts</i>, const struct option <i>*longopts</i>, int <i>*indexptr</i>); int getopt_clip(int <i>argc</i>, char * const <i>argv</i>, const char <i>*shortopts</i>, const struct option <i>*longopts</i>, int <i>*indexptr</i>); extern char <i>*optarg</i>; extern int <i>optind</i>, <i>opterr</i>, <i>optopt</i>;</pre>
DESCRIPTION	<p>These functions are provided as a porting aid for GNU/Freeware/OpenBSD utilities. The <code>getopt_long()</code> function is intended to be as closely compatible with the GNU and OpenBSD implementations as possible, but since these public implementations differ in some corner cases, it is not possible to be fully compatible with both. The differences are enumerated in the NOTES section.</p> <p>The <code>getopt_long()</code> function is an aid for implementing the GNU command line argument conventions. See the GNU documentation for the details of these conventions (glibc 2.2.3). Note that the GNU conventions are not POSIX-conforming. Most notably, the GNU conventions allow for optional option-arguments and do not enforce that operands must follow options on the command line.</p> <p>The <code>getopt_clip()</code> function provides an interface similar to <code>getopt_long()</code> except that it implements the Sun CLIP convention, which is slightly more restrictive than the GNU/Freeware conventions. CLIP is modeled after the GNU/Freeware conventions but removes POSIX violations and syntactic ambiguities (see <code>intro(1)</code>). Specifically, <code>getopt_clip()</code> is a command line parser that can be used by applications that follow the Command Line Interface Paradigm or CLIP syntax guidelines 3, 4, 5, 6, 7, 9, 10, 15, and 16. The remaining guidelines are not addressed by <code>getopt_clip()</code> and are the responsibility of the application.</p> <p>The <code>getopt_long()</code> function is similar to getopt(3C) except that it accepts options in two forms: words and characters, also referred to as long options and short options.</p> <p>The <code>getopt_long()</code> function can be used in two ways. In the first way, every long option understood by the program is mapped to a single character that is usually a corresponding short option. The <code>option</code> structure is used only to translate from long options to short options. In the second way, a long option sets a flag specified in the <code>option</code> structure, or stores a pointer to the command line argument in the address passed to it for options that take arguments. These two methods apply individually to each long option. Both methods can be used in the same application.</p> <p>The <code>getopt_long()</code> function accepts command lines that interleave options and operands. The <code>getopt_long()</code> function reorders the elements of the <code>argv</code> argument such that when all command line arguments have been processed, all operands follow options (and their option-arguments) in the <code>argv</code> array and <code>optind</code> points to the first</p>

operand. The order of options relative to other options and operands relative to other operands is maintained. The argument "--" is accepted as a delimiter indicating the end of options. No argument reorder occurs past this delimiter. Argument reordering can not be unambiguously performed in all cases. The `getopt_long()` function depends on a number of internal heuristics to perform the reordering. The `argc` and `argv` arguments are the argument count and argument array as passed to `main()` (see `exec(2)`).

The `shortopts` argument contains the short-option characters recognized by the command using these functions. If a letter is followed by a colon (:), the option is expected to have an option-argument that should be separated from it by white space. If a character is followed by two colons (::), the option takes an optional option-argument. Any text after the option name it is returned in `optarg`; otherwise, `optarg` is set to 0. A whitespace character can never be used to separate an optional option-argument from its associated option. If `shortopts` contains the character "w" followed by a semicolon (;), then `-w foo` is treated as the long option `--foo`.

If the first character of the `shortopts` argument is the plus sign (+), `getopt_long()` enforces the POSIX requirement that operands follow options on the command line by returning -1 and stopping argument processing upon encountering the first operand (or "--"). This behavior can also be specified by setting the environment variable `POSIXLY_CORRECT`.

A hyphen (-) as the first character of the `shortopts` argument specifies that options and operands can be intermixed in `argv` but no argument reordering is performed. Operands are returned as arguments to option '\1', and option processing does not stop until "--" or the end of `argv` is found.

If the first character of the `shortopts` argument (after a potential plus or minus character) is a colon (:), a colon is returned by `getopt_long()` in response to a missing argument; otherwise, a question mark (?) is returned for this condition.

The `longopts` argument describes the long options to accept. It is an array of `struct option` structures, one for each long option. The array is terminated with an element containing all zeros.

The `struct option` structure contains the following members:

<code>const char *name</code>	Contains a pointer to the name of the option.
<code>int has_arg</code>	Specifies whether the option takes an argument. The possible values, defined in <code><getopt.h></code> , are <code>no_argument</code> , <code>optional_argument</code> , and <code>required_argument</code> .
<code>int *flag</code>	Contains the address of an <code>int</code> variable that is the flag for this option. The value contained in <code>val</code> is stored in this location to indicate that the option was seen. If <code>flag</code> is a null pointer, then the value contained in <code>val</code> is returned when this option is encountered, otherwise zero is returned.

getopt_long(3C)

`int val` Contains the value to be stored at the variable pointed to by `flag` or returned by `getopt_long()` if `flag` is a null pointer.

For any long option, `getopt_long()` returns the index in the array `longopts` of the options definition by storing it in `indexptr`. The name of the option can be retrieved with `longopts[*indexptr].name`. Long options can be distinguished either by the values in their `val` members or by their indices. The `indexptr` variable can also distinguish long options that set flags. The value of `indexptr` after encountering a short option is undefined.

If an option has an argument, the `optarg` global variable is set to point to the start of the option argument on return from `getopt_long()`; otherwise it is set to null. A long option can take an argument in one of two forms: `--option=arg` or `--option arg`. If the long option argument is optional, only the `--option=arg` form can be used to specify the option argument. No argument is specified by the simple form `--option`. The form `--option=` specifies an empty string as the option argument.

Long-option names can be abbreviated if the abbreviation is unique or an exact match for some defined option. An exact match takes precedence over an abbreviated match. Thus, if `foo` and `foobar` are acceptable long-option names, then specifying `--foo` on the command line always matches the former. Specifying `--f` or `--fo` would not be accepted as a match for either.

The `getopt_long()` function places in `optind` the `argv` index of the next argument to be processed. The `optind` global variable is external and is initialized to 1 before the first call to `getopt_long()`. When all options have been processed (that is, up to the first non-option argument), `getopt_long()` returns -1. The special option `--` (two hyphens) can be used to delimit the end of the options; when it is encountered, -1 is returned and `--` is skipped. This option is useful in delimiting non-option arguments that begin with `-` (hyphen).

If `getopt_long()` encounters a short option character `shortopts` string or a long option not described in the `longopts` array, it returns the question mark (?) character. It also returns a question mark (?) character in response to a missing option argument unless the first character of `shortopts` is a colon (:) (or the second character, if the first character is either a plus (+) or a minus (-)), in which case it returns a colon (:). In either case, if the application has not set `opterr` to 0 and the first character of `shortopts` is not a colon (:), `getopt_long()` prints a diagnostic message to `stderr`.

The `getopt_long_only()` function is equivalent to the `getopt_long()` function except that it allows the user of the application to pass long options with only a single hyphen (-) instead of `--`. The `--` prefix is still recognized. However, when a single hyphen (-) is encountered, `getopt_long_only()` attempts to match this argument to a long option, including abbreviations of the long option. If a long option starts with the same character as a short option, a single hyphen followed by that character (and no other characters) will be recognized as a short option. Use of `getopt_long_only()` is strongly discouraged by Sun and GNU for new applications.

The behavior of `getopt_clip()` differs from that of `getopt_long()` in the following ways:

- The `getopt_clip()` function does not perform argument reordering. The `getopt_clip()` function always enforces the POSIX behavior that all options should precede operands on the command line. Specifically, `getopt_clip()` does not reorder arguments but returns -1 and stops processing upon encountering the first operand argument.
- The environment variable `POSIXLY_CORRECT` is ignored (the `getopt_clip()` function behaves as though it were set.)
- The plus and minus characters do not have a special meaning as the first character of the *shortopts* argument. They are treated as any other character (other than the colon) would be treated.
- Optional option-arguments are not allowed. The behavior of `getopt_clip()` when `optional_argument` is specified as the value of `has_arg` in the *longopts* argument or double colons are included in the *shortopts* argument is unspecified.
- Long-option abbreviations are not recognized.
- Short options are required to have at least one long-option equivalent. That is, each character in *shortopts* must appear as the `val` member in one or more `option` structures. Similarly, each long option must have a short option equivalent, meaning that the `val` member of each `option` structure must appear in the *shortopts* string. If these requirements are not met, `getopt_clip()` returns -1 and sets `errno` to `EINVAL`.

RETURN VALUES

For short options (other than `-W` when `W`; is in *shortopts*), these functions return the next option character specified on the command line. For long options, the value returned by these functions depends upon the value of the `flag` structure element for the identified option. If `flag` is `NULL`, the value contained in the `val` structure element for the long option encountered on the command line is returned. Otherwise, these functions return 0 (and the value specified in the `val` member for the long option is stored into the location pointed to by `flag`). When `W`; is in *shortopts* and `-W` is encountered in the command line and the option argument to `-W` matches a long-option name, the return state from these functions is as if the long option had been encountered. However, if no argument is specified to the long option, *optarg* is set to the option argument of `-W` (the long-option name or unique prefix). If the option argument of `-W` does not match a long option (or unique prefix), the return state is as for any other short option.

A colon (:) is returned if `getopt_long()` detects a missing argument and the first character of *shortopts* (other than a possible initial "+" or "-") was a colon (':').

A question mark (?) is returned if `getopt_long()` encounters an option letter not included in *shortopts* or detects a missing argument and the first character of *shortopts* (other than a possible initial "+" or "-") was not a colon (':').

getopt_long(3C)

The `getopt_clip()` function expects all short options to have one or more long-option equivalent and all long options to have one short option equivalent (see NOTES for details). If proper equivalents are not found, `getopt_clip()` returns -1 and sets `errno` to `EINVAL`.

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

`EINVAL` A short option does not have at least one long-option equivalent, or a long option does not have at least one short-option equivalent.

EXAMPLES **EXAMPLE 1** Exmple using `getopt()`.

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

/* Flag set by '--verbose'. */
static int verbose_flag;

int
main (int argc, char **argv)
{
    int c;

    while (1) {
        static struct option long_options[] = {
            /* These options set a flag. */
            {"verbose", no_argument, &verbose_flag, 1},
            {"brief", no_argument, &verbose_flag, 0},
            /* The following options don't set a flag. */
            {"add", no_argument, NULL, 'a'},
            {"append", no_argument, NULL, 'b'},
            {"delete", required_argument, NULL, 'd'},
            {"create", required_argument, NULL, 'c'},
            {"file", required_argument, NULL, 'f'},
            {0, 0, 0, 0}
        };
        /* getopt_long stores the option index here. */
        int option_index = 0;

        c = getopt_long (argc, argv, "abc:d:f:",
                        long_options, &option_index);

        /* Detect the end of the options. */
        if (c == -1)
            break;

        switch (c) {
            case 0:
                /* (In this example) only options which set */
                /* a flag return zero, so do nothing. */
                break;

            case 'a':
                puts ("option --add (-a)\n");
                break;
        }
    }
}
```

EXAMPLE 1 Exmple using getopt (). (Continued)

```

    case 'b':
        puts ("option --append (-b)\n");
        break;

    case 'c':
        printf ("option --create (-c) with value '%s'\n", optarg);
        break;

    case 'd':
        printf ("option --delete (-d) with value '%s'\n", optarg);
        break;

    case 'f':
        printf ("option --file (-f) with value '%s'\n", optarg);
        break;

    case '?':
        /* getopt_long already printed an error message. */
        break;

    default:
        abort ();
}

/* Instead of reporting '--verbose'
and '--brief' as they are encountered,
we report the final status resulting from them. */
if (verbose_flag)
    puts ("verbose flag is set");

/* Print any remaining command line arguments (not options). */
if (optind < argc) {
    printf ("non-option ARGV-elements: ");
    while (optind < argc)
        printf ("%s ", argv[optind++]);
    putchar ('\n');
}

exit (0);
}

```

ENVIRONMENT VARIABLES

See environ(5) for descriptions of the following environment variables that affect the execution of getopt_long(): LANG, LC_ALL, and LC_MESSAGES.

POSIXLY_CORRECT When set (and the first character of the *shoropts* argument is neither a plus or minus sign), the POSIX rule that all operands must follow all options is enforced. Option processing terminates when the first operand is encountered. The getopt_clip() function ignores the setting of POSIXLY_CORRECT and always behaves as if it were set.

getopt_long(3C)

LC_CTYPE Determine the locale for the interpretation of sequences of bytes as characters in *shortopts* and the *longopts*[],*name* structure members.

USAGE The `getopt_long()` function does not fully check for mandatory arguments because there is no unambiguous algorithm to do so. Given an option string `a:b` and the input `-a -b`, `getopt_long()` assumes that `-b` is the mandatory argument to the `-a` option and not that `-a` is missing a mandatory argument. Indeed, the only time a missing option argument can be reliably detected is when the option is the final option on the command line and is not followed by any command arguments.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving
MT-Level	Unsafe

SEE ALSO `intro(1)`, `getopts(1)`, `getopt(3C)`, `getsubopt(3C)`, `gettext(3C)`, `setlocale(3C)`, `attributes(5)`, `environ(5)`, `standards(5)`

NOTES Use of `getopt_long()` is discouraged for applications targeted strictly for Solaris. It should be used only for applications targeted at Solaris and platforms that adhere to the GNU command line conventions. The `getopt_long_only()` function is provided by Solaris and GNU for legacy applications and its use is discouraged by both current conventions.

The differences between the Solaris/GNU and OpenBSD versions of these functions are as follows:

- The handling of the hyphen (-) as the first character of the option string in presence of the environment variable `POSIXLY_CORRECT`:
 - Solaris/GNU Operands are returned as arguments to option `'\1'`, and option processing does not stop until `--` or the end of `argv` is found.
 - OpenBSD obeys `POSIXLY_CORRECT` and stops at the first non-option.
- The handling of the hyphen (-) within the *shortopts* parameter string when not the first character.
 - Solaris/GNU treats a single hyphen (-) on the command line as an operand.
 - OpenBSD treats a single hyphen (-) on the command line as an option. BSD recognizes this behavior as incorrect, but maintains it for compatibility.
- The return value in the event of a missing argument if the first character after `+` or `-` in the option string is not a colon (:)
 - Solaris/GNU returns `?"`.

- OpenBSD returns ":" (since OpenBSD's getopt does).
- The setting *optopt* for long options with *flag* != NULL:
 - Solaris/GNU sets *optopt* to *val*.
 - OpenBSD sets *optopt* to 0 (since *val* would never be returned).
 - The setting of *optarg* for long options without an argument that are invoked with *-w* (*w*; in option string):
 - Solaris/GNU sets *optarg* to the option name (the argument of *-w*).
 - OpenBSD sets *optarg* to NULL (the argument of the long option).
 - The handling of *-w* with an argument that is not (a prefix to) a known long option (*w*; in option string):
 - Solaris/GNU returns '*w*' with *optarg* set to the unknown option.
 - OpenBSD treats as an error (unknown option) and returns "?" with *optopt* set to 0 and *optarg* set to NULL.
 - The error messages are different (all).
 - The implementations do not permute the argument vector at the same points in the calling sequence. The aspects normally used by the caller (ordering after -1 is returned, value of *optind* relative to current positions) are the same. Applications should not depend upon the ordering of the argument vector before -1 is returned.

getpagesize(3C)

NAME	getpagesize – get system page size				
SYNOPSIS	<pre>#include <unistd.h> int getpagesize(void);</pre>				
DESCRIPTION	<p>The <code>getpagesize()</code> function returns the number of bytes in a page. Page granularity is the granularity of many of the memory management calls.</p> <p>The page size is a system page size and need not be the same as the underlying hardware page size.</p> <p>The <code>getpagesize()</code> function is equivalent to <code>sysconf(_SC_PAGE_SIZE)</code> and <code>sysconf(_SC_PAGESIZE)</code>. See sysconf(3C).</p>				
RETURN VALUES	The <code>getpagesize()</code> function returns the current page size.				
ERRORS	No errors are defined.				
USAGE	The value returned by <code>getpagesize()</code> need not be the minimum value that malloc(3C) can allocate. Moreover, the application cannot assume that an object of this size can be allocated with <code>malloc()</code> .				
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>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	pagesize(1) , brk(2) , getrlimit(2) , mmap(2) , mprotect(2) , munmap(2) , malloc(3C) , msync(3C) , sysconf(3C) , attributes(5)				

NAME	getpagesizes – get system supported page sizes				
SYNOPSIS	<pre>#include <sys/mman.h> int getpagesizes(size_t <i>pagesize</i>[], int <i>nelem</i>);</pre>				
DESCRIPTION	The <code>getpagesizes()</code> function returns either the number of different page sizes supported by the system or the actual sizes themselves. When called with <i>nelem</i> as 0 and <i>pagesize</i> as NULL, <code>getpagesizes()</code> returns the number of supported page sizes. Otherwise, up to <i>nelem</i> page sizes are retrieved and assigned to successive elements of <i>pagesize</i> []. The return value is the number of page sizes retrieved and set in <i>pagesize</i> [].				
RETURN VALUES	Upon successful completion, the number of pagesizes supported or actually retrieved is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.				
ERRORS	The <code>getpagesizes()</code> function will fail if: <table border="0" style="margin-left: 2em;"> <tr> <td style="vertical-align: top;">EINVAL</td> <td>The <i>nelem</i> argument is less than 0 or <i>pagesize</i> is NULL but <i>nelem</i> is non-zero.</td> </tr> </table>	EINVAL	The <i>nelem</i> argument is less than 0 or <i>pagesize</i> is NULL but <i>nelem</i> is non-zero.		
EINVAL	The <i>nelem</i> argument is less than 0 or <i>pagesize</i> is NULL but <i>nelem</i> is non-zero.				
USAGE	The <code>getpagesizes()</code> function returns all the page sizes for which the hardware and system software provide support for the <code>mencntl(2)</code> command <code>MC_HAT_ADVISE</code> . Not all processors support all page sizes or combinations of page sizes with equal efficiency. Applications programmers should take this into consideration when using <code>getpagesizes()</code> .				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes: <table border="1" style="margin-left: 2em; 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>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>mencntl(2)</code> , <code>mmap(2)</code> , <code>getpagesize(3C)</code> , <code>attributes(5)</code>				

getpass(3C)

NAME	getpass, getpassphrase – read a string of characters without echo										
SYNOPSIS	<pre>#include <stdlib.h> char *getpass (const char *prompt) ; char *getpassphrase (const char *prompt) ;</pre>										
XPG4, SUS, SUSv2	<pre>#include <unistd.h> char *getpass (const char *prompt) ;</pre>										
DESCRIPTION	<p>The <code>getpass()</code> function opens the process's controlling terminal, writes to that device the null-terminated string <code>prompt</code>, disables echoing, reads a string of characters up to the next newline character or EOF, restores the terminal state and closes the terminal.</p> <p>The <code>getpassphrase()</code> function is identical to <code>getpass()</code>, except that it reads and returns a string of up to 257 characters in length.</p>										
RETURN VALUES	Upon successful completion, <code>getpass()</code> returns a pointer to a null-terminated string of at most 9 bytes that were read from the terminal device. If an error is encountered, the terminal state is restored and a null pointer is returned.										
ERRORS	<p>The <code>getpass()</code> and <code>getpassphrase()</code> functions may fail if:</p> <table><tr><td>EINTR</td><td>The function was interrupted by a signal.</td></tr><tr><td>EIO</td><td>The process is a member of a background process attempting to read from its controlling terminal, the process is ignoring or blocking the SIGTTIN signal or the process group is orphaned.</td></tr><tr><td>EMFILE</td><td>OPEN_MAX file descriptors are currently open in the calling process.</td></tr><tr><td>ENFILE</td><td>The maximum allowable number of files is currently open in the system.</td></tr><tr><td>ENXIO</td><td>The process does not have a controlling terminal.</td></tr></table>	EINTR	The function was interrupted by a signal.	EIO	The process is a member of a background process attempting to read from its controlling terminal, the process is ignoring or blocking the SIGTTIN signal or the process group is orphaned.	EMFILE	OPEN_MAX file descriptors are currently open in the calling process.	ENFILE	The maximum allowable number of files is currently open in the system.	ENXIO	The process does not have a controlling terminal.
EINTR	The function was interrupted by a signal.										
EIO	The process is a member of a background process attempting to read from its controlling terminal, the process is ignoring or blocking the SIGTTIN signal or the process group is orphaned.										
EMFILE	OPEN_MAX file descriptors are currently open in the calling process.										
ENFILE	The maximum allowable number of files is currently open in the system.										
ENXIO	The process does not have a controlling terminal.										
USAGE	The return value points to static data whose content may be overwritten by each call.										
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:										
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td><code>getpass()</code> is Standard.</td></tr><tr><td>MT-Level</td><td>Unsafe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	<code>getpass()</code> is Standard.	MT-Level	Unsafe				
ATTRIBUTE TYPE	ATTRIBUTE VALUE										
Interface Stability	<code>getpass()</code> is Standard.										
MT-Level	Unsafe										
SEE ALSO	<code>attributes(5)</code> , <code>standards(5)</code>										

NAME	getpeerucred – get connected socket/STREAM peer’s credentials														
SYNOPSIS	<pre>#include <ucred.h> int getpeerucred(int <i>fd</i>, ucred_t **<i>ucred</i>);</pre>														
DESCRIPTION	<p>The <code>getpeerucred()</code> function returns the credentials of the peer endpoint of a connection-oriented socket (<code>SOCK_STREAM</code>) or <code>STREAM</code> <i>fd</i> at the time the endpoint was created or the connection was established. A process that initiates a connection retrieves the credentials of its peer at the time the peer’s endpoint was created. A process that listens for connections retrieves the credentials of the peer at the time the peer initiated the connection.</p> <p>When successful, <code>getpeerucred()</code> stores the pointer to a freshly allocated <code>ucred_t</code> in the memory location pointed to by the <code>ucred</code> argument if that memory location contains the null pointer. If the memory location is non-null, it will reuse the existing <code>ucred_t</code>.</p> <p>When <code>ucred</code> is no longer needed, a credential allocated by <code>getpeerucred()</code> should be freed with <code>ucred_free(3C)</code>.</p> <p>It is possible that all fields of the <code>ucred_t</code> are not available to for all peer endpoints and all callers.</p>														
RETURN VALUES	Upon successful completion, <code>getpeerucred()</code> returns 0. Otherwise, it returns <code>-1</code> and <code>errno</code> is set to indicate the error.														
ERRORS	<p>The <code>getpeerucred()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>EAGAIN</code></td> <td>There is not enough memory available to allocate sufficient memory to hold the user credential. The application can try again later.</td> </tr> <tr> <td><code>EBADF</code></td> <td>The <i>fd</i> argument is not a valid file descriptor.</td> </tr> <tr> <td><code>EFAULT</code></td> <td>The pointer location pointed to by the <code>ucred_t **</code> argument points to an invalid, non-null address.</td> </tr> <tr> <td><code>EINVAL</code></td> <td>The socket is connected but the peer credentials are unknown.</td> </tr> <tr> <td><code>ENOMEM</code></td> <td>The physical limits of the system are exceeded by the memory allocation needed to hold the user credential.</td> </tr> <tr> <td><code>ENOTCONN</code></td> <td>The socket or <code>STREAM</code> is not connected or the <code>STREAM</code>’s peer is unknown.</td> </tr> <tr> <td><code>ENOTSUP</code></td> <td>This operation is not supported on this file descriptor.</td> </tr> </table>	<code>EAGAIN</code>	There is not enough memory available to allocate sufficient memory to hold the user credential. The application can try again later.	<code>EBADF</code>	The <i>fd</i> argument is not a valid file descriptor.	<code>EFAULT</code>	The pointer location pointed to by the <code>ucred_t **</code> argument points to an invalid, non-null address.	<code>EINVAL</code>	The socket is connected but the peer credentials are unknown.	<code>ENOMEM</code>	The physical limits of the system are exceeded by the memory allocation needed to hold the user credential.	<code>ENOTCONN</code>	The socket or <code>STREAM</code> is not connected or the <code>STREAM</code> ’s peer is unknown.	<code>ENOTSUP</code>	This operation is not supported on this file descriptor.
<code>EAGAIN</code>	There is not enough memory available to allocate sufficient memory to hold the user credential. The application can try again later.														
<code>EBADF</code>	The <i>fd</i> argument is not a valid file descriptor.														
<code>EFAULT</code>	The pointer location pointed to by the <code>ucred_t **</code> argument points to an invalid, non-null address.														
<code>EINVAL</code>	The socket is connected but the peer credentials are unknown.														
<code>ENOMEM</code>	The physical limits of the system are exceeded by the memory allocation needed to hold the user credential.														
<code>ENOTCONN</code>	The socket or <code>STREAM</code> is not connected or the <code>STREAM</code> ’s peer is unknown.														
<code>ENOTSUP</code>	This operation is not supported on this file descriptor.														
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:														

getpeerucred(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving
MT-Level	MT-Safe

SEE ALSO [door_ucred\(3DOOR\)](#), [ucred_get\(3C\)](#), [attributes\(5\)](#), [connld\(7M\)](#)

NOTES The system currently supports both sides of connection endpoints for local AF_UNIX, AF_INET, and AF_INET6 sockets, /dev/tcp, /dev/ticots, and /dev/ticotsord XTI/TLI connections, and pipe file descriptors sent using I_SENDFD as a result of the open of a named pipe with the "connld" module pushed.

NAME	getpriority, setpriority – get or set process scheduling priority		
SYNOPSIS	<pre>#include <sys/resource.h> int getpriority(int <i>which</i>, id_t <i>who</i>); int setpriority(int <i>which</i>, id_t <i>who</i>, int <i>priority</i>);</pre>		
DESCRIPTION	<p>The <code>getpriority()</code> function obtains the current scheduling priority of a process, process group, or user. The <code>setpriority()</code> function sets the scheduling priority of a process, process group, or user.</p> <p>Target processes are specified by the values of the <i>which</i> and <i>who</i> arguments. The <i>which</i> argument can be one of the following values: <code>PRIO_PROCESS</code>, <code>PRIO_PGRP</code>, <code>PRIO_USER</code>, <code>PRIO_GROUP</code>, <code>PRIO_SESSION</code>, <code>PRIO_LWP</code>, <code>PRIO_TASK</code>, <code>PRIO_PROJECT</code>, <code>PRIO_ZONE</code>, or <code>PRIO_CONTRACT</code>, indicating that the <i>who</i> argument is to be interpreted as a process ID, a process group ID, an effective user ID, an effective group ID, a session ID, an lwp ID, a task ID, a project ID, a zone ID, or a process contract ID, respectively. A 0 value for the <i>who</i> argument specifies the current process, process group, or user. A 0 value for the <i>who</i> argument is treated as valid group ID, session ID, lwp ID, task ID, project ID, zone ID, or process contract ID. A <code>P_MYID</code> value for the <i>who</i> argument can be used to specify the current group, session, lwp, task, project, zone, or process contract, respectively.</p> <p>If more than one process is specified, <code>getpriority()</code> returns the highest priority (lowest numerical value) pertaining to any of the specified processes, and <code>setpriority()</code> sets the priorities of all of the specified processes to the specified value.</p> <p>The default <i>priority</i> is 0; negative priorities cause more favorable scheduling. The range of valid priority values is <code>[-20, 20]</code>. If the value specified to <code>setpriority()</code> is less than the system's lowest supported priority value, the system's lowest supported value is used. If it is greater than the system's highest supported value, the system's highest supported value is used.</p> <p>Only a process with appropriate privileges can raise its priority (that is, assign a lower numerical priority value).</p> <p>Any processes or threads using <code>SCHED_FIFO</code> or <code>SCHED_RR</code> is unaffected by a call to <code>setpriority()</code>. This is not considered an error. A process that subsequently reverts to <code>SCHED_OTHER</code> need not have its priority affected by such a <code>setpriority()</code> call.</p>		
RETURN VALUES	<p>Upon successful completion, <code>getpriority()</code> returns an integer in the range from <code>-20</code> to <code>20</code>. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.</p> <p>Upon successful completion, <code>setpriority()</code> returns <code>0</code>. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.</p>		
ERRORS	<p>The <code>getpriority()</code> and <code>setpriority()</code> functions will fail if:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><code>ESRCH</code></td> <td>No process could be located using the <i>which</i> and <i>who</i> argument values specified.</td> </tr> </table>	<code>ESRCH</code>	No process could be located using the <i>which</i> and <i>who</i> argument values specified.
<code>ESRCH</code>	No process could be located using the <i>which</i> and <i>who</i> argument values specified.		

getpriority(3C)

EINVAL The value of the *which* argument was not recognized, or the value of the *who* argument is not a valid process ID, process group ID, user ID, group ID, session ID, lwp ID, task ID, project ID, or zone ID.

In addition, `setpriority()` may fail if:

EPERM A process was located, but neither the real nor effective user ID of the executing process is the privileged user or match the effective user ID of the process whose priority is being changed.

EACCES A request was made to change the priority to a lower numeric value (that is, to a higher priority) and the current process does not have appropriate privileges.

EXAMPLES

EXAMPLE 1 Example using `getpriority()`

The following example returns the current scheduling priority for the process ID returned by the call to `getpid(2)`.

```
#include <sys/resource.h>
...
int which = PRIO_PROCESS;
id_t pid;
int ret;

pid = getpid();
ret = getpriority(which, pid);
```

EXAMPLE 2 Example using `setpriority()`

The following example sets the priority for the current process ID to -20.

```
#include <sys/resource.h>
...
int which = PRIO_PROCESS;
id_t pid;
int priority = -20;
int ret;

pid = getpid();
ret = setpriority(which, pid, priority);
```

USAGE

The effect of changing the scheduling priority can vary depending on the process-scheduling algorithm in effect.

Because `getpriority()` can return -1 on successful completion, it is necessary to set `errno` to 0 prior to a call to `getpriority()`. If `getpriority()` returns -1, then `errno` can be checked to see if an error occurred or if the value is a legitimate priority.

getpriority(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

SEE ALSO `nice(1)`, `renice(1)`, `fork(2)`, `sched_get_priority_max(3RT)`,
`sched_setscheduler(3RT)`, `attributes(5)`, `standards(5)`

getpw(3C)

NAME	getpw – get passwd entry from UID				
SYNOPSIS	<pre>#include <stdlib.h> int getpw(uid_t <i>uid</i>, char *<i>buf</i>);</pre>				
DESCRIPTION	The <code>getpw()</code> function searches the user data base for a user id number that equals <i>uid</i> , copies the line of the password file in which <i>uid</i> was found into the array pointed to by <i>buf</i> , and returns 0. <code>getpw()</code> returns non-zero if <i>uid</i> cannot be found.				
USAGE	<p>This function is included only for compatibility with prior systems and should not be used; the functions described on the getpwnam(3C) manual page should be used instead.</p> <p>If the <code>/etc/passwd</code> and the <code>/etc/group</code> files have a plus sign (+) for the NIS entry, then <code>getpwent()</code> and <code>getgrent()</code> will not return NULL when the end of file is reached. See getpwnam(3C).</p>				
RETURN VALUES	The <code>getpw()</code> function returns non-zero on error.				
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>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe				
SEE ALSO	getpwnam(3C) , passwd(4) , attributes(5)				

NAME	getpwnam, getpwnam_r, getpwent, getpwent_r, getpwuid, getpwuid_r, setpwent, endpwent, fgetpwent, fgetpwent_r – get password entry
SYNOPSIS	<pre>#include <pwd.h> struct passwd *getpwnam(const char *name); struct passwd *getpwnam_r(const char *name, struct passwd *pwd, char *buffer, int buflen); struct passwd *getpwent(void); struct passwd *getpwent_r(struct passwd *pwd, char *buffer, int buflen); struct passwd *getpwuid(uid_t uid); struct passwd *getpwuid_r(uid_t uid, struct passwd *pwd, char *buffer, int buflen); void setpwent(void); void endpwent(void); struct passwd *fgetpwent(FILE *f); struct passwd *fgetpwent_r(FILE *f, struct passwd *pwd, char *buffer, int buflen);</pre>
Standard conforming	<pre>cc [flag...] file... -D_POSIX_PTHREAD_SEMANTICS [library...] int getpwnam_r(const char *name, struct passwd *pwd, char *buffer, size_t bufsize, struct passwd **result); int getpwuid_r(uid_t uid, struct passwd *pwd, char *buffer, size_t bufsize, struct passwd **result);</pre>
DESCRIPTION	<p>These functions are used to obtain password entries. Entries can come from any of the sources for passwd specified in the <code>/etc/nsswitch.conf</code> file (see <code>nsswitch.conf(4)</code>).</p> <p>The <code>getpwnam()</code> function searches for a password entry with the login name specified by the character string parameter <i>name</i>.</p> <p>The <code>getpwuid()</code> function searches for a password entry with the (numeric) user ID specified by the <i>uid</i> parameter.</p> <p>The <code>setpwent()</code>, <code>getpwent()</code>, and <code>endpwent()</code> functions are used to enumerate password entries from the database. The <code>setpwent()</code> function sets (or resets) the enumeration to the beginning of the set of password entries. This function should be called before the first call to <code>getpwent()</code>. Calls to <code>getpwnam()</code> and <code>getpwuid()</code> leave the enumeration position in an indeterminate state. Successive calls to <code>getpwent()</code> return either successive entries or a null pointer, indicating the end of the enumeration.</p>

getpwnam(3C)

The `endpwent()` function may be called to indicate that the caller expects to do no further password retrieval operations; the system may then close the password file, deallocate resources it was using, and so forth. It is still allowed, but possibly less efficient, for the process to call more password functions after calling `endpwent()`.

The `fgetpwent()` function, unlike the other functions above, does not use `nsswitch.conf` but reads and parses the next line from the stream `f`, which is assumed to have the format of the `passwd` file. See `passwd(4)`.

Reentrant Interfaces

The `getpwnam()`, `getpwuid()`, `getpwent()`, and `fgetpwent()` functions use thread-specific data storage that is reused in each call to one of these functions by the same thread, making them safe to use but not recommended for multithreaded applications.

The parallel functions `getpwnam_r()`, `getpwuid_r()`, `getpwent_r()`, and `fgetpwent_r()` provide reentrant interfaces for these operations.

Each reentrant interface performs the same operation as its non-reentrant counterpart, named by removing the “`_r`” suffix. The reentrant interfaces, however, use buffers supplied by the caller to store returned results instead of using thread-specific data that can be overwritten by each call. They are safe for use in both single-threaded and multithreaded applications.

Each reentrant interface takes the same parameters as its non-reentrant counterpart, as well as the following additional parameters. The `pwd` parameter must be a pointer to a `struct passwd` structure allocated by the caller. On successful completion, the function returns the password entry in this structure. The parameter `buffer` is a pointer to a buffer supplied by the caller, used as storage space for the password data. All pointers within the returned `struct passwd` `pwd` point to data stored within this buffer; see `passwd Structure` below. The buffer must be large enough to hold all the data associated with the password entry. The parameter `buflen` (or `bufsize` for the standard-conforming versions; see `standards(5)`) should give the size in bytes of `buffer`. The maximum size needed for this buffer can be determined with the `{_SC_GETPW_R_SIZE_MAX} sysconf(3C)` parameter. The standard-conforming versions place a pointer to the modified `pwd` structure in the `result` parameter, instead of returning a pointer to this structure. A null pointer is returned at the location pointed to by `result` on error or if the requested entry is not found.

For enumeration in multithreaded applications, the position within the enumeration is a process-wide property shared by all threads. The `setpwent()` function can be used in a multithreaded application but resets the enumeration position for all threads. If multiple threads interleave calls to `getpwent_r()`, the threads will enumerate disjoint subsets of the password database.

Like their non-reentrant counterparts, `getpwnam_r()` and `getpwuid_r()` leave the enumeration position in an indeterminate state.

passwd Structure

Password entries are represented by the `struct passwd` structure defined in `<pwd.h>`:


```

struct passwd {
    char *pw_name;          /* user's login name */
    char *pw_passwd;       /* no longer used */
    uid_t pw_uid;          /* user's uid */
    gid_t pw_gid;          /* user's gid */
    char *pw_age;          /* not used */
    char *pw_comment;      /* not used */
    char *pw_gecos;        /* typically user's full name */
    char *pw_dir;          /* user's home dir */
    char *pw_shell;        /* user's login shell */
};

```

The `pw_passwd` member should not be used as the encrypted password for the user; use `getspnam()` or `getspnam_r()` instead. See [getspnam\(3C\)](#).

RETURN VALUES

The `getpwnam()`, `getpwnam_r()`, `getpwuid()`, and `getpwuid_r()` functions each return a pointer to a `struct passwd` if they successfully locate the requested entry. A null pointer is returned if the requested entry is not found, or an error occurs. On error, `errno` is set to indicate the error. Upon successful completion (including the case when the requested entry is not found), the standard-conforming functions `getpwnam_r()` and `getpwuid_r()` return 0. Otherwise, an error number is returned to indicate the error.

The `getpwent()`, `getpwent_r()`, `fgetpwent()`, and `fgetpwent_r()` functions each return a pointer to a `struct passwd` if they successfully enumerate an entry; otherwise they return a null pointer on end-of-file or error. On error, `errno` is set to indicate the error.

See [Intro\(2\)](#) for the proper usage and interpretation of `errno` in multithreaded applications.

The `getpwnam()`, `getpwuid()`, `getpwent()`, and `fgetpwent()` functions use thread-specific data storage, so returned data must be copied before a subsequent call to any of these functions if the data is to be saved.

When the pointer returned by the reentrant functions `getpwnam_r()`, `getpwuid_r()`, `getpwent_r()`, and `fgetpwent_r()` is non-null, it is always equal to the `pwd` pointer that was supplied by the caller.

Applications wishing to check for error situations should set `errno` to 0 before calling `getpwnam()`, `getpwnam_r()`, `getpwuid()`, `getpwuid_r()`, `getpwent()`, `getpwent_r()`, `fgetpwent()`, and `fgetpwent_r()`. If these functions return a null pointer and `errno` is non-zero, an error occurred.

ERRORS

The `getpwent_r()`, `fgetpwent()`, and `fgetpwent_r()` functions will fail if:

- `EIO` An I/O error has occurred.
- `ERANGE` Insufficient storage was supplied by `buffer` and `bufsize` to contain the data to be referenced by the resulting `passwd` structure.

The `getpwent_r()` function will fail if:

- `EMFILE` There are `{OPEN_MAX}` file descriptors currently open in the calling process.

getpwnam(3C)

ENFILE The maximum allowable number of files is currently open in the system.

The `getpwnam()`, `getpwnam_r()`, `getpwuid()`, `getpwuid_r()`, `getpwent()`, `setpwent()`, and `endpwent()` functions may fail if:

EIO An I/O error has occurred.

The `getpwnam()`, `getpwnam_r()`, `getpwuid()`, `getpwuid_r()`, `getpwent()`, and `setpwent()` functions may fail if:

EMFILE There are `{OPEN_MAX}` file descriptors currently open in the calling process.

ENFILE The maximum allowable number of files is currently open in the system.

The `getpwnam()`, `getpwnam_r()`, `getpwuid()`, and `getpwuid_r()` functions may fail if:

EINTR A signal was caught during the execution of the function call.

The `getpwnam_r()` and `getpwuid_r()` functions may fail if:

ERANGE Insufficient storage was supplied by *buffer* and *bufsize* to contain the data to be referenced by the resulting `passwd` structure.

USAGE Three names associated with the current process can be determined: `getpwuid(geteuid())` returns the name associated with the effective user ID of the process; `getlogin()` returns the name associated with the current login activity; and `getpwuid(getuid())` returns the name associated with the real user ID of the process.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	See below.
MT-Level	See Reentrant Interfaces in DESCRIPTION.

The `endpwent()`, `getpwent()`, `getpwnam()`, `getpwnam_r()`, `getpwuid()`, `getpwuid_r()`, and `setpwent()` functions are Standard.

SEE ALSO `nispasswd(1)`, `passwd(1)`, `yppasswd(1)`, `Intro(2)`, `Intro(3)`, `cuserid(3C)`, `getgrnam(3C)`, `getlogin(3C)`, `getspnam(3C)`, `nsswitch.conf(4)`, `passwd(4)`, `shadow(4)`, `attributes(5)`, `standards(5)`

NOTES When compiling multithreaded programs, see `Intro(3)`.

Use of the enumeration interfaces `getpwent()` and `getpwent_r()` is discouraged; enumeration is supported for the `passwd` file, NIS, and NIS+, but in general is not efficient and might not be supported for all database sources. The semantics of enumeration are discussed further in `nsswitch.conf(4)`.

Previous releases allowed the use of '+' and '-' entries in `/etc/passwd` to selectively include and exclude NIS entries. The primary usage of these '+/-' entries is superseded by the name service switch, so the '+/-' form might not be supported in future releases.

If required, the '+/-' functionality can still be obtained for NIS by specifying `compat` as the source for `passwd`.

If the '+/-' functionality is required in conjunction with NIS+, specify both `compat` as the source for `passwd` and `nisplus` as the source for the pseudo-database `passwd_compat`. See `passwd(4)`, `shadow(4)`, and `nsswitch.conf(4)` for details.

If the '+/-' is used, both `/etc/shadow` and `/etc/passwd` should have the same '+' and '-' entries to ensure consistency between the password and shadow databases.

If a password entry from any of the sources contains an empty `uid` or `gid` field, that entry will be ignored by the files, NIS, and NIS+ name service switch backends, causing the user to appear unknown to the system.

If a password entry contains an empty `gecos`, `home directory`, or `shell` field, `getpwnam()` and `getpwnam_r()` return a pointer to a null string in the respective field of the `passwd` structure.

If the shell field is empty, `login(1)` automatically assigns the default shell. See `login(1)`.

Solaris 2.4 and earlier releases provided definitions of the `getpwnam_r()` and `getpwuid_r()` functions as specified in POSIX.1c Draft 6. The final POSIX.1c standard changed the interface for these functions. Support for the Draft 6 interface is provided for compatibility only and might not be supported in future releases. New applications and libraries should use the standard-conforming interface.

For POSIX.1c-conforming applications, the `_POSIX_PTHREAD_SEMANTICS` and `_REENTRANT` flags are automatically turned on by defining the `_POSIX_C_SOURCE` flag with a value $\geq 199506L$.

getrusage(3C)

NAME	getrusage – get information about resource utilization										
SYNOPSIS	<pre>#include <sys/resource.h> int getrusage(int who, struct rusage *r_usage);</pre>										
DESCRIPTION	<p>The <code>getrusage()</code> function provides measures of the resources used by the current process, its terminated and waited-for child processes, or the current light weight process (LWP). If the value of the <i>who</i> argument is <code>RUSAGE_SELF</code>, information is returned about resources used by the current process. If the value of the <i>who</i> argument is <code>RUSAGE_CHILDREN</code>, information is returned about resources used by the terminated and waited-for children of the current process. If the child is never waited for (for instance, if the parent has <code>SA_NOCLDWAIT</code> set or sets <code>SIGCHLD</code> to <code>SIG_IGN</code>), the resource information for the child process is discarded and not included in the resource information provided by <code>getrusage()</code>. If the value of the <i>who</i> argument is <code>RUSAGE_LWP</code>, information is returned about resources used by the current LWP.</p> <p>The <i>r_usage</i> argument is a pointer to an object of type <code>struct rusage</code> in which the returned information is stored. The members of <code>rusage</code> are as follows:</p> <pre>struct timeval ru_utime; /* user time used */ struct timeval ru_stime; /* system time used */ long ru_maxrss; /* maximum resident set size */ long ru_idrss; /* integral resident set size */ long ru_minflt; /* page faults not requiring physical I/O */ long ru_majflt; /* page faults requiring physical I/O */ long ru_nswap; /* swaps */ long ru_inblock; /* block input operations */ long ru_oublock; /* block output operations */ long ru_msgsnd; /* messages sent */ long ru_msgrcv; /* messages received */ long ru_nsignals; /* signals received */ long ru_nvcsw; /* voluntary context switches */ long ru_nivcsw; /* involuntary context switches */</pre> <p>The structure members are interpreted as follows:</p> <table><tr><td><code>ru_utime</code></td><td>The total amount of time spent executing in user mode. Time is given in seconds and microseconds.</td></tr><tr><td><code>ru_stime</code></td><td>The total amount of time spent executing in system mode. Time is given in seconds and microseconds.</td></tr><tr><td><code>ru_maxrss</code></td><td>The maximum resident set size. Size is given in pages (the size of a page, in bytes, is given by the <code>getpagesize(3C)</code> function). See the NOTES section of this page.</td></tr><tr><td><code>ru_idrss</code></td><td>An “integral” value indicating the amount of memory in use by a process while the process is running. This value is the sum of the resident set sizes of the process running when a clock tick occurs. The value is given in pages times clock ticks. It does not take sharing into account. See the NOTES section of this page.</td></tr><tr><td><code>ru_minflt</code></td><td>The number of page faults serviced which did not require any physical I/O activity. See the NOTES section of this page.</td></tr></table>	<code>ru_utime</code>	The total amount of time spent executing in user mode. Time is given in seconds and microseconds.	<code>ru_stime</code>	The total amount of time spent executing in system mode. Time is given in seconds and microseconds.	<code>ru_maxrss</code>	The maximum resident set size. Size is given in pages (the size of a page, in bytes, is given by the <code>getpagesize(3C)</code> function). See the NOTES section of this page.	<code>ru_idrss</code>	An “integral” value indicating the amount of memory in use by a process while the process is running. This value is the sum of the resident set sizes of the process running when a clock tick occurs. The value is given in pages times clock ticks. It does not take sharing into account. See the NOTES section of this page.	<code>ru_minflt</code>	The number of page faults serviced which did not require any physical I/O activity. See the NOTES section of this page.
<code>ru_utime</code>	The total amount of time spent executing in user mode. Time is given in seconds and microseconds.										
<code>ru_stime</code>	The total amount of time spent executing in system mode. Time is given in seconds and microseconds.										
<code>ru_maxrss</code>	The maximum resident set size. Size is given in pages (the size of a page, in bytes, is given by the <code>getpagesize(3C)</code> function). See the NOTES section of this page.										
<code>ru_idrss</code>	An “integral” value indicating the amount of memory in use by a process while the process is running. This value is the sum of the resident set sizes of the process running when a clock tick occurs. The value is given in pages times clock ticks. It does not take sharing into account. See the NOTES section of this page.										
<code>ru_minflt</code>	The number of page faults serviced which did not require any physical I/O activity. See the NOTES section of this page.										

<code>ru_majflt</code>	The number of page faults serviced which required physical I/O activity. This could include page ahead operations by the kernel. See the NOTES section of this page.
<code>ru_nswap</code>	The number of times a process was swapped out of main memory.
<code>ru_inblock</code>	The number of times the file system had to perform input in servicing a <code>read(2)</code> request.
<code>ru_oublock</code>	The number of times the file system had to perform output in servicing a <code>write(2)</code> request.
<code>ru_msgsnd</code>	The number of messages sent over sockets.
<code>ru_msgrcv</code>	The number of messages received from sockets.
<code>ru_nsignals</code>	The number of signals delivered.
<code>ru_nvcsw</code>	The number of times a context switch resulted due to a process voluntarily giving up the processor before its time slice was completed (usually to await availability of a resource).
<code>ru_nivcsw</code>	The number of times a context switch resulted due to a higher priority process becoming runnable or because the current process exceeded its time slice.

RETURN VALUES Upon successful completion, `getrusage()` returns 0. Otherwise, `-1` is returned and `errno` is set to indicate the error.

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

<code>EFAULT</code>	The address specified by the <code>r_usage</code> argument is not in a valid portion of the process' address space.
<code>EINVAL</code>	The <code>who</code> parameter is not a valid value.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

SEE ALSO `sar(1M)`, `read(2)`, `times(2)`, `write(2)`, `getpagesize(3C)`, `gettimeofday(3C)`, `wait(3C)`, `attributes(5)`, `standards(5)`

NOTES The `ru_maxrss`, `ru_ixrss`, `ru_idrss`, and `ru_isrss` members of the `rusage` structure are set to 0 in this implementation.

The numbers `ru_inblock` and `ru_oublock` account only for real I/O, and are approximate measures at best. Data supplied by the cache mechanism is charged only to the first process to read and the last process to write the data.

getrusage(3C)

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

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

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

By definition, major page faults require physical I/O, while minor page faults do not require physical I/O. For example, reclaiming the page from the free list would avoid I/O and generate a minor page fault. More commonly, minor page faults occur during process startup as references to pages which are already in memory. For example, if an address space faults on some "hot" executable or shared library, this results in a minor page fault for the address space. Also, any one doing a `read(2)` or `write(2)` to something that is in the page cache will get a minor page fault(s) as well.

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

NAME	gets, fgets – get a string from a stream						
SYNOPSIS	<pre>#include <stdio.h> char *gets(char *s); char *fgets(char *s, int n, FILE *stream);</pre>						
DESCRIPTION	<p>The <code>gets()</code> function reads bytes from the standard input stream (see <code>intro(3)</code>), <code>stdin</code>, into the array pointed to by <code>s</code>, until a newline character is read or an end-of-file condition is encountered. The newline character is discarded and the string is terminated with a null byte.</p> <p>If the length of an input line exceeds the size of <code>s</code>, indeterminate behavior may result. For this reason, it is strongly recommended that <code>gets()</code> be avoided in favor of <code>fgets()</code>.</p> <p>The <code>fgets()</code> function reads bytes from the <code>stream</code> into the array pointed to by <code>s</code>, until <code>n-1</code> bytes are read, or a newline character is read and transferred to <code>s</code>, or an end-of-file condition is encountered. The string is then terminated with a null byte.</p> <p>The <code>fgets()</code> and <code>gets()</code> functions may mark the <code>st_atime</code> field of the file associated with <code>stream</code> for update. The <code>st_atime</code> field will be marked for update by the first successful execution of <code>fgetc(3C)</code>, <code>fgets()</code>, <code>fread(3C)</code>, <code>fscanf(3C)</code>, <code>getc(3C)</code>, <code>getchar(3C)</code>, <code>gets()</code>, or <code>scanf(3C)</code> using <code>stream</code> that returns data not supplied by a prior call to <code>ungetc(3C)</code> or <code>ungetwc(3C)</code>.</p>						
RETURN VALUES	<p>If end-of-file is encountered and no bytes have been read, no bytes are transferred to <code>s</code> and a null pointer is returned. For standard-conforming (see <code>standards(5)</code>) applications, if the end-of-file indicator for the stream is set, no bytes are transferred to <code>s</code> and a null pointer is returned whether or not the stream is at end-of-file. If a read error occurs, such as trying to use these functions on a file that has not been opened for reading, a null pointer is returned and the error indicator for the stream is set. If end-of-file is encountered, the EOF indicator for the stream is set. Otherwise <code>s</code> is returned.</p>						
ERRORS	Refer to <code>fgetc(3C)</code> .						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>lseek(2)</code> , <code>read(2)</code> , <code>ferror(3C)</code> , <code>fgetc(3C)</code> , <code>fgetwc(3C)</code> , <code>fopen(3C)</code> , <code>fread(3C)</code> , <code>getchar(3C)</code> , <code>scanf(3C)</code> , <code>stdio(3C)</code> , <code>ungetc(3C)</code> , <code>ungetwc(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

getspnam(3C)

NAME	getspnam, getspnam_r, getspent, getspent_r, setspent, endspent, fgetspent, fgetspent_r – get password entry
SYNOPSIS	<pre>#include <shadow.h> struct spwd *getspnam(const char *name); struct spwd *getspnam_r(const char *name, struct spwd *result, char *buffer, int buflen); struct spwd *getspent(void); struct spwd *getspent_r(struct spwd *result, char *buffer, int buflen); void setspent(void); void endspent(void); struct spwd *fgetspent(FILE *fp); struct spwd *fgetspent_r(FILE *fp, struct spwd *result, char *buffer, int buflen);</pre>
DESCRIPTION	<p>These functions are used to obtain shadow password entries. An entry may come from any of the sources for shadow specified in the <code>/etc/nsswitch.conf</code> file (see <code>nsswitch.conf(4)</code>).</p> <p>The <code>getspnam()</code> function searches for a shadow password entry with the login name specified by the character string argument <i>name</i>.</p> <p>The <code>setspent()</code>, <code>getspent()</code>, and <code>endspent()</code> functions are used to enumerate shadow password entries from the database.</p> <p>The <code>setspent()</code> function sets (or resets) the enumeration to the beginning of the set of shadow password entries. This function should be called before the first call to <code>getspent()</code>. Calls to <code>getspnam()</code> leave the enumeration position in an indeterminate state.</p> <p>Successive calls to <code>getspent()</code> return either successive entries or NULL, indicating the end of the enumeration.</p> <p>The <code>endspent()</code> function may be called to indicate that the caller expects to do no further shadow password retrieval operations; the system may then close the shadow password file, deallocate resources it was using, and so forth. It is still allowed, but possibly less efficient, for the process to call more shadow password functions after calling <code>endspent()</code>.</p> <p>The <code>fgetspent()</code> function, unlike the other functions above, does not use <code>nsswitch.conf</code>; it reads and parses the next line from the stream <i>fp</i>, which is assumed to have the format of the shadow file (see <code>shadow(4)</code>).</p>
Reentrant Interfaces	The <code>getspnam()</code> , <code>getspent()</code> , and <code>fgetspent()</code> functions use thread-specific data storage that is reused in each call to one of these functions by the same thread, making them safe to use but not recommended for multithreaded applications.

The `getspnam_r()`, `getspent_r()`, and `fgetspent_r()` functions provide reentrant interfaces for these operations.

Each reentrant interface performs the same operation as its non-reentrant counterpart, named by removing the `_r` suffix. The reentrant interfaces, however, use buffers supplied by the caller to store returned results, and are safe for use in both single-threaded and multithreaded applications.

Each reentrant interface takes the same argument as its non-reentrant counterpart, as well as the following additional arguments. The *result* argument must be a pointer to a `struct spwd` structure allocated by the caller. On successful completion, the function returns the shadow password entry in this structure. The *buffer* argument must be a pointer to a buffer supplied by the caller. This buffer is used as storage space for the shadow password data. All of the pointers within the returned `struct spwd result` point to data stored within this buffer (see RETURN VALUES). The buffer must be large enough to hold all of the data associated with the shadow password entry. The *buflen* argument should give the size in bytes of the buffer indicated by *buffer*.

For enumeration in multithreaded applications, the position within the enumeration is a process-wide property shared by all threads. The `setspent()` function may be used in a multithreaded application but resets the enumeration position for all threads. If multiple threads interleave calls to `getspent_r()`, the threads will enumerate disjoint subsets of the shadow password database.

Like its non-reentrant counterpart, `getspnam_r()` leaves the enumeration position in an indeterminate state.

RETURN VALUES

Password entries are represented by the `struct spwd` structure defined in `<shadow.h>`:

```
struct spwd{
    char          *sp_namp;      /* login name */
    char          *sp_pwdp;     /* encrypted passwd */
    long          sp_lstchg;     /* date of last change */
    long          sp_min;       /* min days to passwd change */
    long          sp_max;       /* max days to passwd change*/
    long          sp_warn;      /* warning period */
    long          sp_inact;     /* max days inactive */
    long          sp_expire;    /* account expiry date */
    unsigned long sp_flag;      /* not used */
};
```

See `shadow(4)` for more information on the interpretation of this data.

The `getspnam()` and `getspnam_r()` functions each return a pointer to a `struct spwd` if they successfully locate the requested entry; otherwise they return `NULL`.

The `getspent()`, `getspent_r()`, `fgetspent()`, and `fgetspent_r()` functions each return a pointer to a `struct spwd` if they successfully enumerate an entry; otherwise they return `NULL`, indicating the end of the enumeration.

getspnam(3C)

The `getspnam()`, `getspent()`, and `fgetspent()` functions use thread-specific data storage, so returned data must be copied before a subsequent call to any of these functions if the data is to be saved.

When the pointer returned by the reentrant functions `getspnam_r()`, `getspent_r()`, and `fgetspent_r()` is non-null, it is always equal to the *result* pointer that was supplied by the caller.

ERRORS The reentrant functions `getspnam_r()`, `getspent_r()`, and `fgetspent_r()` will return NULL and set `errno` to `ERANGE` if the length of the buffer supplied by caller is not large enough to store the result. See `intro(2)` for the proper usage and interpretation of `errno` in multithreaded applications.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	See "Reentrant Interfaces" in <code>DESCRIPTION</code> .

SEE ALSO `nispasswd(1)`, `passwd(1)`, `yppasswd(1)`, `intro(3)`, `getlogin(3C)`, `getpwnam(3C)`, `nsswitch.conf(4)`, `passwd(4)`, `shadow(4)`, `attributes(5)`

WARNINGS The reentrant interfaces `getspnam_r()`, `getspent_r()`, and `fgetspent_r()` are included in this release on an uncommitted basis only, and are subject to change or removal in future minor releases.

NOTES When compiling multithreaded applications, see `intro(3)`, *Notes On Multithreaded Applications*, for information about the use of the `_REENTRANT` flag.

Use of the enumeration interfaces `getspent()` and `getspent_r()` is not recommended; enumeration is supported for the shadow file, NIS, and NIS+, but in general is not efficient and may not be supported for all database sources. The semantics of enumeration are discussed further in `nsswitch.conf(4)`.

Access to shadow password information may be restricted in a manner depending on the database source being used. Access to the `/etc/shadow` file is generally restricted to processes running with the effective uid of the file owner of the `{PRIV_FILE_DAC_READ}` privilege. Other database sources may impose stronger or less stringent restrictions.

When NIS is used as the database source, the information for the shadow password entries is obtained from the "passwd.byname" map. This map stores only the information for the `sp_namp` and `sp_pwdp` fields of the `struct spwd` structure. Shadow password entries obtained from NIS will contain the value `-1` in the remainder of the fields.

getspnam(3C)

When NIS+ is used as the database source, and the caller lacks the permission needed to retrieve the encrypted password from the NIS+ "passwd.org_dir" table, the NIS+ service returns the string "*NP*" instead of the actual encrypted password string. The functions described on this page will then return the string "*NP*" to the caller as the value of the member `sp_pwdp` in the returned shadow password structure.

getsubopt(3C)

NAME	getsubopt – parse suboptions from a string
SYNOPSIS	<pre>#include <stdlib.h> int getsubopt (char **<i>optionp</i>, char * const *<i>tokens</i>, char **<i>valuep</i>);</pre>
DESCRIPTION	<p>The <code>getsubopt()</code> function parses suboptions in a flag argument that was initially parsed by <code>getopt(3C)</code>. The suboptions are separated by commas and may consist of either a single token or a token-value pair separated by an equal sign. Since commas delimit suboptions in the option string, they are not allowed to be part of the suboption or the value of a suboption; if present in the option input string, they are changed to null characters. White-space characters within tokens or token-value pairs must be protected from the shell by quotes.</p> <p>The syntax described above is used in the following example by the <code>mount(1M)</code>, utility, which allows the user to specify mount parameters with the <code>-o</code> option as follows:</p> <pre>% mount -o rw,hard,bg,wsiz=1024 speed:/usr /usr</pre> <p>In this example there are four suboptions: <code>rw</code>, <code>hard</code>, <code>bg</code>, and <code>wsiz</code>, the last of which has an associated value of 1024.</p> <p>The <code>getsubopt()</code> function takes the address of a pointer to the option string, a vector of possible tokens, and the address of a value string pointer. It returns the index of the token that matched the suboption in the input string, or <code>-1</code> if there was no match. If the option string pointed to by <code>*optionp</code> contains only one suboption, <code>getsubopt()</code> updates <code>*optionp</code> to point to the null character at the end of the string; otherwise it isolates the suboption by replacing the comma separator with a null character, and updates <code>*optionp</code> to point to the start of the next suboption. If the suboption has an associated value, <code>getsubopt()</code> updates <code>*valuep</code> to point to the value's first character. Otherwise it sets <code>*valuep</code> to <code>NULL</code>.</p> <p>The token vector is organized as a series of pointers to null strings. The end of the token vector is identified by a null pointer.</p> <p>When <code>getsubopt()</code> returns, a non-null value for <code>*valuep</code> indicates that the suboption that was processed included a value. The calling program may use this information to determine if the presence or absence of a value for this suboption is an error.</p> <p>When <code>getsubopt()</code> fails to match the suboption with the tokens in the <code>tokens</code> array, the calling program should decide if this is an error, or if the unrecognized option should be passed to another program.</p>
RETURN VALUES	<p>The <code>getsubopt()</code> function returns the index of the matched token string, or <code>-1</code> when the token it is scanning is not in the token vector. The variable addressed by <code>*valuep</code> contains a pointer to the first character of the token that was not recognized, rather than a pointer to a value for that token.</p> <p>The variable addressed by <code>*optionp</code> points to the next option to be parsed, or a null character if there are no more options.</p>

ERRORS No errors are defined.

EXAMPLES **EXAMPLE 1** Use `getsubopt()` to process options.

The following example demonstrates the processing of options to the `mount(1M)` utility using `getsubopt()`.

```
#include <stdlib.h>

char *myopts[] = {
#define READONLY      0
    "ro",
#define READWRITE     1
    "rw",
#define WRITESIZE     2
    "wsize",
#define READSIZE      3
    "rsize",
    NULL};

main(argc, argv)
    int argc;
    char **argv;
{
    int sc, c, errflag;
    char *options, *value;
    extern char *optarg;
    extern int optind;
    .
    .
    .
    while((c = getopt(argc, argv, "abf:o:")) != -1) {
        switch (c) {
            case 'a': /* process a option */
                break;
            case 'b': /* process b option */
                break;
            case 'f':
                ofile = optarg;
                break;
            case '?':
                errflag++;
                break;
            case 'o':
                options = optarg;
                while (*options != '\0') {
                    switch(getsubopt(&options, myopts, &value)) {
                        case READONLY : /* process ro option */
                            break;
                        case READWRITE : /* process rw option */
                            break;
                        case WRITESIZE : /* process wsize option */
                            if (value == NULL) {
                                error_no_arg();
                                errflag++;
                            } else

```

getsubopt(3C)

EXAMPLE 1 Use `getsubopt()` to process options. (Continued)

```
        write_size = atoi(value);
        break;
    case READSIZE : /* process rsize option */
        if (value == NULL) {
            error_no_arg();
            errflag++;
        } else
            read_size = atoi(value);
        break;
    default :
        /* process unknown token */
        error_bad_token(value);
        errflag++;
        break;
    }
}
break;
}
}
if (errflag) {
    /* print usage instructions etc. */
}
for (; optind < argc; optind++) {
    /* process remaining arguments */
}
.
.
}
```

EXAMPLE 2 Parse suboptions.

The following example uses the `getsubopt()` function to parse a value argument in the `optarg` external variable returned by a call to `getopt(3C)`.

```
#include <stdlib.h>
...
char *tokens[] = {"HOME", "PATH", "LOGNAME", (char *) NULL };
char *value;
int opt, index;
while ((opt = getopt(argc, argv, "e:")) != -1) {
    switch(opt) {
        case 'e' :
            while ((index = getsubopt(&optarg, tokens, &value)) != -1) {
                switch(index) {
                    ...
                }
            }
            break;
        ...
    }
}
```

getsubopt(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `mount(1M)`, `getopt(3C)`, `attributes(5)`, `standards(5)`

gettext(3C)

NAME	gettext, dgettext, dcgettext, ngettext, dngettext, dcngettext, textdomain, bindtextdomain, bind_textdomain_codeset – message handling functions
SYNOPSIS	
Solaris and GNU-compatible	<pre>#include <libintl.h> char *gettext(const char *msgid); char *dgettext(const char *domainname, const char *msgid); char *textdomain(const char *domainname); char *bindtextdomain(const char *domainname, const char *dirname); #include <libintl.h> #include <locale.h> char *dcgettext(const char *domainname, const char *msgid, int category);</pre>
GNU-compatible	<pre>#include <libintl.h> char *ngettext(const char *msgid1, const char *msgid2, unsigned long int n); char *dngettext(const char *domainname, const char *msgid1, const char *msgid2, unsigned long int n); char *bind_textdomain_codeset(const char *domainname, const char *codeset); #include <libintl.h> #include <locale.h> char *dcngettext(const char *domainname, const char *msgid1, const char *msgid2, unsigned long int n, int category);</pre>
DESCRIPTION	<p>The <code>gettext()</code>, <code>dgettext()</code>, and <code>dcgettext()</code> functions attempt to retrieve a target string based on the specified <code>msgid</code> argument within the context of a specific domain and the current locale. The length of strings returned by <code>gettext()</code>, <code>dgettext()</code>, and <code>dcgettext()</code> is undetermined until the function is called. The <code>msgid</code> argument is a null-terminated string.</p> <p>The <code>ngettext()</code>, <code>dngettext()</code>, and <code>dcngettext()</code> functions are equivalent to <code>gettext()</code>, <code>dgettext()</code>, and <code>dcgettext()</code>, respectively, except for the handling of plural forms. These functions work only with GNU-compatible message catalogues. The <code>ngettext()</code>, <code>dngettext()</code>, and <code>dcngettext()</code> functions search for the message string using the <code>msgid1</code> argument as the key and the <code>n</code> argument to determine the plural form. If no message catalogues are found, <code>msgid1</code> is returned if <code>n == 1</code>, otherwise <code>msgid2</code> is returned.</p> <p>The <code>NLSPATH</code> environment variable (see <code>environ(5)</code>) is searched first for the location of the <code>LC_MESSAGES</code> catalogue. The setting of the <code>LC_MESSAGES</code> category of the current locale determines the locale used by <code>gettext()</code> and <code>dgettext()</code> for string</p>

retrieval. The *category* argument determines the locale used by `dcgettext()`. If `NLSPATH` is not defined and the current locale is "C", `gettext()`, `dgettext()`, and `dcgettext()` simply return the message string that was passed. In a locale other than "C", if `NLSPATH` is not defined or if a message catalogue is not found in any of the components specified by `NLSPATH`, the routines search for the message catalogue using the scheme described in the following paragraph.

The `LANGUAGE` environment variable is examined to determine the GNU-compatible message catalogues to be used. The value of `LANGUAGE` is a list of locale names separated by a colon (':') character. If `LANGUAGE` is defined, each locale name is tried in the specified order and if a GNU-compatible message catalogue is found, the message is returned. If a GNU-compatible message catalogue is found but failed to find a corresponding *msgid*, the *msgid* string is return. If `LANGUAGE` is not defined or if a Solaris message catalogue is found or no GNU-compatible message catalogue is found in processing `LANGUAGE`, the pathname used to locate the message catalogue is *dirname/locale/category/domainname.mo*, where *dirname* is the directory specified by `bindtextdomain()`, *locale* is a locale name, and *category* is either `LC_MESSAGES` if `gettext()`, `dgettext()`, `ngettext()`, or `dngettext()` is called, or `LC_XXX` where the name is the same as the locale category name specified by the *category* argument to `dcgettext()` or `dcngettext()`.

For `gettext()` and `ngettext()`, the domain used is set by the last valid call to `textdomain()`. If a valid call to `textdomain()` has not been made, the default domain (called `messages`) is used.

For `dgettext()`, `dcgettext()`, `dngettext()`, and `dcngettext()`, the domain used is specified by the *domainname* argument. The *domainname* argument is equivalent in syntax and meaning to the *domainname* argument to `textdomain()`, except that the selection of the domain is valid only for the duration of the `dgettext()`, `dcgettext()`, `dngettext()`, or `dcngettext()` function call.

The `textdomain()` function sets or queries the name of the current domain of the active `LC_MESSAGES` locale category. The *domainname* argument is a null-terminated string that can contain only the characters allowed in legal filenames.

The *domainname* argument is the unique name of a domain on the system. If there are multiple versions of the same domain on one system, namespace collisions can be avoided by using `bindtextdomain()`. If `textdomain()` is not called, a default domain is selected. The setting of domain made by the last valid call to `textdomain()` remains valid across subsequent calls to `setlocale(3C)`, and `gettext()`.

The *domainname* argument is applied to the currently active `LC_MESSAGES` locale.

The current setting of the domain can be queried without affecting the current state of the domain by calling `textdomain()` with *domainname* set to the null pointer. Calling `textdomain()` with a *domainname* argument of a null string sets the domain to the default domain (`messages`).

gettext(3C)

The `bindtextdomain()` function binds the path predicate for a message domain *domainname* to the value contained in *dirname*. If *domainname* is a non-empty string and has not been bound previously, `bindtextdomain()` binds *domainname* with *dirname*.

If *domainname* is a non-empty string and has been bound previously, `bindtextdomain()` replaces the old binding with *dirname*. The *dirname* argument can be an absolute or relative pathname being resolved when `gettext()`, `dgettext()`, or `dcgettext()` are called. If *domainname* is a null pointer or an empty string, `bindtextdomain()` returns NULL. User defined domain names cannot begin with the string `SYS_`. Domain names beginning with this string are reserved for system use.

The `bind_textdomain_codeset()` function can be used to specify the output codeset for message catalogues for domain *domainname*. The *codeset* argument must be a valid codeset name that can be used for the `iconv_open(3C)` function, or a null pointer. If the *codeset* argument is the null pointer, `bind_textdomain_codeset()` returns the currently selected codeset for the domain with the name *domainname*. It returns a null pointer if a codeset has not yet been selected. The `bind_textdomain_codeset()` function can be used multiple times. If used multiple times with the same *domainname* argument, the later call overrides the settings made by the earlier one. The `bind_textdomain_codeset()` function returns a pointer to a string containing the name of the selected codeset. The string is allocated internally in the function and must not be changed by the user.

RETURN VALUES

The `gettext()`, `dgettext()`, and `dcgettext()` functions return the message string if the search succeeds. Otherwise they return the *msgid* string.

The `ngettext()`, `dngettext()`, and `dcngettext()` functions return the message string if the search succeeds. If the search fails, *msgid1* is returned if *n* == 1. Otherwise *msgid2* is returned.

The individual bytes of the string returned by `gettext()`, `dgettext()`, `dcgettext()`, `ngettext()`, `dngettext()`, or `dcngettext()` can contain any value other than NULL. If *msgid* is a null pointer, the return value is undefined. The string returned must not be modified by the program and can be invalidated by a subsequent call to `bind_textdomain_codeset()` or `setlocale(3C)`. If the *domainname* argument to `dgettext()`, `dcgettext()`, `dngettext()`, or `dcngettext()` is a null pointer, the the domain currently bound by `textdomain()` is used.

The normal return value from `textdomain()` is a pointer to a string containing the current setting of the domain. If *domainname* is a null pointer, `textdomain()` returns a pointer to the string containing the current domain. If `textdomain()` was not previously called and *domainname* is a null string, the name of the default domain is returned. The name of the default domain is `messages`. If `textdomain()` fails, a null pointer is returned.

The return value from `bindtextdomain()` is a null-terminated string containing *dirname* or the directory binding associated with *domainname* if *dirname* is NULL. If no binding is found, the default return value is `/usr/lib/locale`. If *domainname* is a

null pointer or an empty string, `bindtextdomain()` takes no action and returns a null pointer. The string returned must not be modified by the caller. If `bindtextdomain()` fails, a null pointer is returned.

USAGE These functions impose no limit on message length. However, a text *domainname* is limited to `TEXTDOMAINMAX` (256) bytes.

The `gettext()`, `dgettext()`, `dcgettext()`, `ngettext()`, `dngettext()`, `dcngettext()`, `textdomain()`, and `bindtextdomain()` functions can be used safely in multithreaded applications, as long as `setlocale(3C)` is not being called to change the locale.

The `gettext()`, `dgettext()`, `dcgettext()`, `textdomain()`, and `bindtextdomain()` functions work with both Solaris message catalogues and GNU-compatible message catalogues. The `ngettext()`, `dngettext()`, `dcngettext()`, and `bind_textdomain_codeset()` functions work only with GNU-compatible message catalogues. See `msgfmt(1)` for information about Solaris message catalogues and GNU-compatible message catalogues.

FILES `/usr/lib/locale`
 default path predicate for message domain files

`/usr/lib/locale/locale/LC_MESSAGES/domainname.mo`
 system default location for file containing messages for language *locale* and *domainname*

`/usr/lib/locale/locale/LC_XXX/domainname.mo`
 system default location for file containing messages for language *locale* and *domainname* for `dcgettext()` calls where `LC_XXX` is `LC_CTYPE`, `LC_NUMERIC`, `LC_TIME`, `LC_COLLATE`, `LC_MONETARY`, or `LC_MESSAGES`

`dirname/locale/LC_MESSAGES/domainname.mo`
 location for file containing messages for domain *domainname* and path predicate *dirname* after a successful call to `bindtextdomain()`

`dirname/locale/LC_XXX/domainname.mo`
 location for files containing messages for domain *domainname*, language *locale*, and path predicate *dirname* after a successful call to `bindtextdomain()` for `dcgettext()` calls where `LC_XXX` is one of `LC_CTYPE`, `LC_NUMERIC`, `LC_TIME`, `LC_COLLATE`, `LC_MONETARY`, or `LC_MESSAGES`

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe with exceptions

SEE ALSO `msgfmt(1)`, `xgettext(1)`, `iconv_open(3C)`, `setlocale(3C)`, `attributes(5)`, `environ(5)`

gettimeofday(3C)

NAME	gettimeofday, settimeofday – get or set the date and time						
SYNOPSIS	<pre>#include <sys/time.h> int gettimeofday(struct timeval *tp, void *tzp); int settimeofday(struct timeval *tp, void *tzp);</pre>						
DESCRIPTION	<p>The <code>gettimeofday()</code> function gets and the <code>settimeofday()</code> function sets the system's notion of the current time. The current time is expressed in elapsed seconds and microseconds since 00:00 Universal Coordinated Time, January 1, 1970. The resolution of the system clock is hardware dependent; the time may be updated continuously or in clock ticks.</p> <p>The <code>tp</code> argument points to a <code>timeval</code> structure, which includes the following members:</p> <pre>long tv_sec; /* seconds since Jan. 1, 1970 */ long tv_usec; /* and microseconds */</pre> <p>If <code>tp</code> is a null pointer, the current time information is not returned or set.</p> <p>The TZ environment variable holds time zone information. See TIMEZONE(4).</p> <p>The <code>tzp</code> argument to <code>gettimeofday()</code> and <code>settimeofday()</code> is ignored.</p> <p>Only privileged processes can set the time of day.</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 <code>settimeofday()</code> function will fail if:</p> <table><tr><td>EINVAL</td><td>The structure pointed to by <code>tp</code> specifies an invalid time.</td></tr><tr><td>EPERM</td><td>The {PRIV_SYS_TIME} privilege was not asserted in the effective set of the calling process.</td></tr></table> <p>The <code>gettimeofday()</code> function will fail for 32-bit interfaces if:</p> <table><tr><td>E_OVERFLOW</td><td>The system time has progressed beyond 2038, thus the size of the <code>tv_sec</code> member of the <code>timeval</code> structure pointed to by <code>tp</code> is insufficient to hold the current time in seconds.</td></tr></table>	EINVAL	The structure pointed to by <code>tp</code> specifies an invalid time.	EPERM	The {PRIV_SYS_TIME} privilege was not asserted in the effective set of the calling process.	E_OVERFLOW	The system time has progressed beyond 2038, thus the size of the <code>tv_sec</code> member of the <code>timeval</code> structure pointed to by <code>tp</code> is insufficient to hold the current time in seconds.
EINVAL	The structure pointed to by <code>tp</code> specifies an invalid time.						
EPERM	The {PRIV_SYS_TIME} privilege was not asserted in the effective set of the calling process.						
E_OVERFLOW	The system time has progressed beyond 2038, thus the size of the <code>tv_sec</code> member of the <code>timeval</code> structure pointed to by <code>tp</code> is insufficient to hold the current time in seconds.						
USAGE	If the <code>tv_usec</code> member of <code>tp</code> is > 500000, <code>settimeofday()</code> rounds the seconds upward. If the time needs to be set with better than one second accuracy, call <code>settimeofday()</code> for the seconds and then <code>adjtime(2)</code> for finer accuracy.						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:						

gettimeofday(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	gettimeofday() is Standard.
MT-Level	MT-Safe

SEE ALSO adjtime(2), ctime(3C), gethrtime(3C), TIMEZONE(4), attributes(5), privileges(5), standards(5)

gettimeofday(3UCB)

NAME	gettimeofday, settimeofday – get or set the date and time
SYNOPSIS	<pre><code>/usr/ucb/cc[flag ...] file ... #include <sys/time.h> int gettimeofday(tp, tzp); struct timeval *tzp; struct timezone *tzp; int settimeofday(tp, tzp); struct timeval *tzp; struct timezone *tzp;</code></pre>
DESCRIPTION	<p>The system's notion of the current Greenwich time is obtained with the <code>gettimeofday()</code> call, and set with the <code>settimeofday()</code> call. The current time is expressed in elapsed seconds and microseconds since 00:00 GMT, January 1, 1970 (zero hour). The resolution of the system clock is hardware dependent; the time may be updated continuously, or in clock ticks.</p> <pre><code>long tv_sec; /* seconds since Jan. 1, 1970 */ long tv_usec; /* and microseconds */</code></pre> <p><i>tp</i> points to a <code>timeval</code> structure, which includes the following members:</p> <p>If <i>tp</i> is a NULL pointer, the current time information is not returned or set.</p> <p><i>tzp</i> is an obsolete pointer formerly used to get and set timezone information. <i>tzp</i> is now ignored. Timezone information is now handled using the TZ environment variable; see <code>TIMEZONE(4)</code>.</p> <p>Only the privileged user may set the time of day.</p>
RETURN VALUES	A -1 return value indicates an error occurred; in this case an error code is stored in the global variable <code>errno</code> .
ERRORS	The following error codes may be set in <code>errno</code> : EINVAL <i>tp</i> specifies an invalid time. EPERM A user other than the privileged user attempted to set the time.
SEE ALSO	<code>adjtime(2)</code> , <code>ctime(3C)</code> , <code>gettimeofday(3C)</code> , <code>TIMEZONE(4)</code>
NOTES	Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported. <i>tzp</i> is ignored in SunOS 5.X releases. <code>tv_usec</code> is always 0.

NAME	gettxt – retrieve a text string
SYNOPSIS	<pre>#include <nl_types.h> char *gettxt(const char *msgid, const char *dflt_str);</pre>
DESCRIPTION	<p>The <code>gettxt()</code> function retrieves a text string from a message file. The arguments to the function are a message identification <code>msgid</code> and a default string <code>dflt_str</code> to be used if the retrieval fails.</p> <p>The text strings are in files created by the <code>mkmsgs</code> utility (see <code>mkmsgs(1)</code>) and installed in directories in <code>/usr/lib/locale/locale/LC_MESSAGES</code>.</p> <p>The directory <code>locale</code> can be viewed as the language in which the text strings are written. The user can request that messages be displayed in a specific language by setting the environment variable <code>LC_MESSAGES</code>. If <code>LC_MESSAGES</code> is not set, the environment variable <code>LANG</code> will be used. If <code>LANG</code> is not set, the files containing the strings are in <code>/usr/lib/locale/C/LC_MESSAGES/*</code>.</p> <p>The user can also change the language in which the messages are displayed by invoking the <code>setlocale(3C)</code> function with the appropriate arguments.</p> <p>If <code>gettxt()</code> fails to retrieve a message in a specific language it will try to retrieve the same message in U.S. English. On failure, the processing depends on what the second argument <code>dflt_str</code> points to. A pointer to the second argument is returned if the second argument is not the null string. If <code>dflt_str</code> points to the null string, a pointer to the U.S. English text string "Message not found!!\n" is returned.</p> <p>The following depicts the acceptable syntax of <code>msgid</code> for a call to <code>gettxt()</code>.</p> <pre><msgid> = <msgfilename>:<msgnumber></pre> <p>The first field is used to indicate the file that contains the text strings and must be limited to 14 characters. These characters must be selected from the set of all character values excluding <code>\0</code> (null) and the ASCII code for <code>/</code> (slash) and <code>:</code> (colon). The names of message files must be the same as the names of files created by <code>mkmsgs</code> and installed in <code>/usr/lib/locale/locale/LC_MESSAGES/*</code>. The numeric field indicates the sequence number of the string in the file. The strings are numbered from 1 to <code>n</code> where <code>n</code> is the number of strings in the file.</p>
RETURN VALUES	Upon failure to pass either the correct <code>msgid</code> or a valid message number to <code>gettxt()</code> , a pointer to the text string "Message not found!!\n" is returned.
USAGE	It is recommended that <code>gettext(3C)</code> be used in place of this function.
EXAMPLES	<p>EXAMPLE 1 Example of <code>gettxt()</code> function.</p> <p>In the following example,</p>

gettext(3C)

EXAMPLE 1 Example of `gettext()` function. (Continued)

```
gettext("UX:10", "hello world\n")
gettext("UX:10", "")
```

UX is the name of the file that contains the messages and 10 is the message number.

FILES

`/usr/lib/locale/C/LC_MESSAGES/*`

contains default message files created by `mkmsgs`

`/usr/lib/locale/locale/LC_MESSAGES/*`

contains message files for different languages created by `mkmsgs`

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe with exceptions

SEE ALSO

`exstr(1)`, `mkmsgs(1)`, `schtxt(1)`, `gettext(3C)`, `fmtmsg(3C)`, `setlocale(3C)`, `attributes(5)`, `environ(5)`

NAME	getusershell, setusershell, endusershell – get legal user shells																								
SYNOPSIS	<pre>#include <unistd.h> char *getusershell(void); void setusershell(void); void endusershell(void);</pre>																								
DESCRIPTION	<p>The <code>getusershell()</code> function returns a pointer to a legal user shell as defined by the system manager in the file <code>/etc/shells</code>. If <code>/etc/shells</code> does not exist, the following locations of the standard system shells are used in its place:</p> <table border="0"> <tr><td><code>/bin/bash</code></td><td><code>/bin/csh</code></td></tr> <tr><td><code>/bin/jsh</code></td><td><code>/bin/ksh</code></td></tr> <tr><td><code>/bin/pfcsh</code></td><td><code>/bin/pfksh</code></td></tr> <tr><td><code>/bin/pfsh</code></td><td><code>/bin/sh</code></td></tr> <tr><td><code>/bin/tcsh</code></td><td><code>/bin/zsh</code></td></tr> <tr><td><code>/sbin/jsh</code></td><td><code>/sbin/pfsh</code></td></tr> <tr><td><code>/sbin/sh</code></td><td><code>/usr/bin/bash</code></td></tr> <tr><td><code>/usr/bin/csh</code></td><td><code>/usr/bin/jsh</code></td></tr> <tr><td><code>/usr/bin/ksh</code></td><td><code>/usr/bin/pfcsh</code></td></tr> <tr><td><code>/usr/bin/pfksh</code></td><td><code>/usr/bin/pfsh</code></td></tr> <tr><td><code>/usr/bin/sh</code></td><td><code>/usr/bin/tcsh</code></td></tr> <tr><td><code>/usr/bin/zsh</code></td><td><code>/usr/xpg4/bin/sh</code></td></tr> </table> <p>The <code>getusershell()</code> function opens the file <code>/etc/shells</code>, if it exists, and returns the next entry in the list of shells.</p> <p>The <code>setusershell()</code> function rewinds the file or the list.</p> <p>The <code>endusershell()</code> function closes the file, frees any memory used by <code>getusershell()</code> and <code>setusershell()</code>, and rewinds the file <code>/etc/shells</code>.</p>	<code>/bin/bash</code>	<code>/bin/csh</code>	<code>/bin/jsh</code>	<code>/bin/ksh</code>	<code>/bin/pfcsh</code>	<code>/bin/pfksh</code>	<code>/bin/pfsh</code>	<code>/bin/sh</code>	<code>/bin/tcsh</code>	<code>/bin/zsh</code>	<code>/sbin/jsh</code>	<code>/sbin/pfsh</code>	<code>/sbin/sh</code>	<code>/usr/bin/bash</code>	<code>/usr/bin/csh</code>	<code>/usr/bin/jsh</code>	<code>/usr/bin/ksh</code>	<code>/usr/bin/pfcsh</code>	<code>/usr/bin/pfksh</code>	<code>/usr/bin/pfsh</code>	<code>/usr/bin/sh</code>	<code>/usr/bin/tcsh</code>	<code>/usr/bin/zsh</code>	<code>/usr/xpg4/bin/sh</code>
<code>/bin/bash</code>	<code>/bin/csh</code>																								
<code>/bin/jsh</code>	<code>/bin/ksh</code>																								
<code>/bin/pfcsh</code>	<code>/bin/pfksh</code>																								
<code>/bin/pfsh</code>	<code>/bin/sh</code>																								
<code>/bin/tcsh</code>	<code>/bin/zsh</code>																								
<code>/sbin/jsh</code>	<code>/sbin/pfsh</code>																								
<code>/sbin/sh</code>	<code>/usr/bin/bash</code>																								
<code>/usr/bin/csh</code>	<code>/usr/bin/jsh</code>																								
<code>/usr/bin/ksh</code>	<code>/usr/bin/pfcsh</code>																								
<code>/usr/bin/pfksh</code>	<code>/usr/bin/pfsh</code>																								
<code>/usr/bin/sh</code>	<code>/usr/bin/tcsh</code>																								
<code>/usr/bin/zsh</code>	<code>/usr/xpg4/bin/sh</code>																								
RETURN VALUES	The <code>getusershell()</code> function returns a null pointer on EOF.																								
BUGS	All information is contained in memory that may be freed with a call to <code>endusershell()</code> , so it must be copied if it is to be saved.																								

getutent(3C)

NAME	getutent, getutid, getutline, pututline, setutent, endutent, utmpname – user accounting database functions
SYNOPSIS	<pre>#include <utmp.h> struct utmp *getutent(void); struct utmp *getutid(const struct utmp *id); struct utmp *getutline(const struct utmp *line); struct utmp *pututline(const struct utmp *utmp); void setutent(void); void endutent(void); int utmpname(const char *file);</pre>
DESCRIPTION	<p>These functions provide access to the user accounting database, utmp. Entries in the database are described by the definitions and data structures in <utmp.h>.</p> <p>The utmp structure contains the following members:</p> <pre>char ut_user[8]; /* user login name */ char ut_id[4]; /* /sbin/inittab id (usually line #) */ char ut_line[12]; /* device name (console, lnxx) */ short ut_pid; /* process id */ short ut_type; /* type of entry */ struct exit_status ut_exit; /* exit status of a process */ /* marked as DEAD_PROCESS */ time_t ut_time; /* time entry was made */</pre> <p>The structure exit_status includes the following members:</p> <pre>short e_termination; /* termination status */ short e_exit; /* exit status */</pre>
getutent()	The getutent() function reads in the next entry from a utmp database. If the database is not already open, it opens it. If it reaches the end of the database, it fails.
getutid()	The getutid() function searches forward from the current point in the utmp database until it finds an entry with a ut_type matching id->ut_type if the type specified is RUN_LVL, BOOT_TIME, DOWN_TIME, OLD_TIME, or NEW_TIME. If the type specified in id is INIT_PROCESS, LOGIN_PROCESS, USER_PROCESS, or DEAD_PROCESS, then getutid() will return a pointer to the first entry whose type is one of these four and whose ut_id member matches id->ut_id. If the end of database is reached without a match, it fails.
getutline()	The getutline() function searches forward from the current point in the utmp database until it finds an entry of the type LOGIN_PROCESS or ut_line string matching the line->ut_line string. If the end of database is reached without a match, it fails.

pututline()	The <code>pututline()</code> function writes the supplied <code>utmp</code> structure into the <code>utmp</code> database. It uses <code>getutid()</code> to search forward for the proper place if it finds that it is not already at the proper place. It is expected that normally the user of <code>pututline()</code> will have searched for the proper entry using one of the these functions. If so, <code>pututline()</code> will not search. If <code>pututline()</code> does not find a matching slot for the new entry, it will add a new entry to the end of the database. It returns a pointer to the <code>utmp</code> structure. When called by a non-root user, <code>pututline()</code> invokes a <code>setuid()</code> root program to verify and write the entry, since the <code>utmp</code> database is normally writable only by root. In this event, the <code>ut_name</code> member must correspond to the actual user name associated with the process; the <code>ut_type</code> member must be either <code>USER_PROCESS</code> or <code>DEAD_PROCESS</code> ; and the <code>ut_line</code> member must be a device special file and be writable by the user.				
setutent()	The <code>setutent()</code> function resets the input stream to the beginning. This reset should be done before each search for a new entry if it is desired that the entire database be examined.				
endutent()	The <code>endutent()</code> function closes the currently open database.				
utmpname()	The <code>utmpname()</code> function allows the user to change the name of the database file examined to another file. If the file does not exist, this will not be apparent until the first attempt to reference the file is made. The <code>utmpname()</code> function does not open the file but closes the old file if it is currently open and saves the new file name.				
RETURN VALUES	A null pointer is returned upon failure to read, whether for permissions or having reached the end of file, or upon failure to write. If the file name given is longer than 79 characters, <code>utmpname()</code> returns 0. Otherwise, it returns 1.				
USAGE	<p>These functions use buffered standard I/O for input, but <code>pututline()</code> uses an unbuffered non-standard write to avoid race conditions between processes trying to modify the <code>utmp</code> and <code>wtmp</code> databases.</p> <p>Applications should not access the <code>utmp</code> and <code>wtmp</code> databases directly, but should use these functions to ensure that these databases are maintained consistently. Using these functions, however, may cause applications to fail if user accounting data cannot be represented properly in the <code>utmp</code> structure (for example, on a system where PIDs can exceed 32767). Use the functions described on the getutxent(3C) manual page instead.</p>				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<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>Unsafe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Unsafe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Unsafe				
SEE ALSO	getutxent(3C) , ttyslot(3C) , utmpx(4) , attributes(5)				
NOTES	The most current entry is saved in a static structure. Multiple accesses require that it be copied before further accesses are made. On each call to either <code>getutid()</code> or <code>getutline()</code> , the function examines the static structure before performing more I/O.				

getutent(3C)

If the contents of the static structure match what it is searching for, it looks no further. For this reason, to use `getutline()` to search for multiple occurrences, it would be necessary to zero out the static area after each success, or `getutline()` would just return the same structure over and over again. There is one exception to the rule about emptying the structure before further reads are done. The implicit read done by `pututline()` (if it finds that it is not already at the correct place in the file) will not hurt the contents of the static structure returned by the `getutent()`, `getutid()` or `getutline()` functions, if the user has just modified those contents and passed the pointer back to `pututline()`.

NAME	getutxent, getutxid, getutxline, pututxline, setutxent, endutxent, utmpxname, getutmp, getutmpx, updwtmp, updwtmpx – user accounting database functions
SYNOPSIS	<pre>#include <utmpx.h> struct utmpx *getutxent(void); struct utmpx *getutxid(const struct utmpx *id); struct utmpx *getutxline(const struct utmpx *line); struct utmpx *pututxline(const struct utmpx *utmpx); void setutxent(void); void endutxent(void); int utmpxname(const char *file); void getutmp(struct utmpx *utmpx, struct utmp *utmp); void getutmpx(struct utmp *utmp, struct utmpx *utmpx); void updwtmp(char *wfile, struct utmp *utmp); void updwtmpx(char *wfilex, struct utmpx *utmpx);</pre>
DESCRIPTION	<p>These functions provide access to the user accounting database, utmpx (see utmpx(4)). Entries in the database are described by the definitions and data structures in <utmpx.h>.</p> <p>The utmpx structure contains the following members:</p> <pre>char ut_user[32]; /* user login name */ char ut_id[4]; /* /etc/inittab id (usually line #) */ char ut_line[32]; /* device name (console, lnxx) */ pid_t ut_pid; /* process id */ short ut_type; /* type of entry */ struct exit_status ut_exit; /* exit status of a process */ /* marked as DEAD_PROCESS */ struct timeval ut_tv; /* time entry was made */ int ut_session; /* session ID, used for windowing */ short ut_syslen; /* significant length of ut_host */ /* including terminating null */ char ut_host[257]; /* host name, if remote */</pre> <p>The exit_status structure includes the following members:</p> <pre>short e_termination; /* termination status */ short e_exit; /* exit status */</pre> <p>getutxent() The getutxent() function reads in the next entry from a utmpx database. If the database is not already open, it opens it. If it reaches the end of the database, it fails.</p> <p>getutxid() The getutxid() function searches forward from the current point in the utmpx database until it finds an entry with a ut_type matching id->ut_type, if the type specified is RUN_LVL, BOOT_TIME, DOWN_TIME, OLD_TIME, or NEW_TIME. If the type specified in id is INIT_PROCESS, LOGIN_PROCESS, USER_PROCESS, or</p>

getutxent(3C)

	DEAD_PROCESS, then <code>getutxid()</code> will return a pointer to the first entry whose type is one of these four and whose <code>ut_id</code> member matches <code>id->ut_id</code> . If the end of database is reached without a match, it fails.
<code>getutxline()</code>	The <code>getutxline()</code> function searches forward from the current point in the <code>utmpx</code> database until it finds an entry of the type <code>LOGIN_PROCESS</code> or <code>USER_PROCESS</code> which also has a <code>ut_line</code> string matching the <code>line->ut_line</code> string. If the end of the database is reached without a match, it fails.
<code>pututxline()</code>	The <code>pututxline()</code> function writes the supplied <code>utmpx</code> structure into the <code>utmpx</code> database. It uses <code>getutxid()</code> to search forward for the proper place if it finds that it is not already at the proper place. It is expected that normally the user of <code>pututxline()</code> will have searched for the proper entry using one of the <code>getutx()</code> routines. If so, <code>pututxline()</code> will not search. If <code>pututxline()</code> does not find a matching slot for the new entry, it will add a new entry to the end of the database. It returns a pointer to the <code>utmpx</code> structure. When called by a non-root user, <code>pututxline()</code> invokes a <code>setuid()</code> root program to verify and write the entry, since the <code>utmpx</code> database is normally writable only by root. In this event, the <code>ut_name</code> member must correspond to the actual user name associated with the process; the <code>ut_type</code> member must be either <code>USER_PROCESS</code> or <code>DEAD_PROCESS</code> ; and the <code>ut_line</code> member must be a device special file and be writable by the user.
<code>setutxent()</code>	The <code>setutxent()</code> function resets the input stream to the beginning. This should be done before each search for a new entry if it is desired that the entire database be examined.
<code>endutxent()</code>	The <code>endutxent()</code> function closes the currently open database.
<code>utmpxname()</code>	The <code>utmpxname()</code> function allows the user to change the name of the database file examined from <code>/var/adm/utmpx</code> to any other file, most often <code>/var/adm/wtmpx</code> . If the file does not exist, this will not be apparent until the first attempt to reference the file is made. The <code>utmpxname()</code> function does not open the file, but closes the old file if it is currently open and saves the new file name. The new file name must end with the "x" character to allow the name of the corresponding <code>utmp</code> file to be easily obtainable.; otherwise, an error value of 0 is returned. The function returns 1 on success.
<code>getutmp()</code>	The <code>getutmp()</code> function copies the information stored in the members of the <code>utmpx</code> structure to the corresponding members of the <code>utmp</code> structure. If the information in any member of <code>utmpx</code> does not fit in the corresponding <code>utmp</code> member, the data is silently truncated. (See getutent(3C) for <code>utmp</code> structure)
<code>getutmpx()</code>	The <code>getutmpx()</code> function copies the information stored in the members of the <code>utmp</code> structure to the corresponding members of the <code>utmpx</code> structure. (See getutent(3C) for <code>utmp</code> structure)
<code>updwtmp()</code>	The <code>updwtmp()</code> function can be used in two ways.

	<p>If <i>wfile</i> is <code>/var/adm/wtmp</code>, the <code>utmp</code> format record supplied by the caller is converted to a <code>utmpx</code> format record and the <code>/var/adm/wtmpx</code> file is updated (because the <code>/var/adm/wtmp</code> file no longer exists, operations on <code>wtmp</code> are converted to operations on <code>wtmpx</code> by the library functions).</p> <p>If <i>wfile</i> is a file other than <code>/var/adm/wtmp</code>, it is assumed to be an old file in <code>utmp</code> format and is updated directly with the <code>utmp</code> format record supplied by the caller.</p>
updwtmpx ()	<p>The <code>updwtmpx ()</code> function writes the contents of the <code>utmpx</code> structure pointed to by <i>utmpx</i> to the database.</p>
utmpx structure	<p>The values of the <code>e_termination</code> and <code>e_exit</code> members of the <code>ut_exit</code> structure are valid only for records of type <code>DEAD_PROCESS</code>. For <code>utmpx</code> entries created by <code>init(1M)</code>, these values are set according to the result of the <code>wait ()</code> call that <code>init</code> performs on the process when the process exits. See the wait(3C), manual page for the values <code>init</code> uses. Applications creating <code>utmpx</code> entries can set <code>ut_exit</code> values using the following code example:</p> <pre>u->ut_exit.e_termination = WTERMSIG(process->p_exit) u->ut_exit.e_exit = WEXITSTATUS(process->p_exit)</pre> <p>See <code>wait.h(3HEAD)</code> for descriptions of the <code>WTERMSIG</code> and <code>WEXITSTATUS</code> macros.</p> <p>The <code>ut_session</code> member is not acted upon by the operating system. It is used by applications interested in creating <code>utmpx</code> entries.</p> <p>For records of type <code>USER_PROCESS</code>, the <code>nonuser ()</code> and <code>nonuserx ()</code> macros use the value of the <code>ut_exit.e_exit</code> member to mark <code>utmpx</code> entries as real logins (as opposed to multiple <code>xterms</code> started by the same user on a window system). This allows the system utilities that display users to obtain an accurate indication of the number of actual users, while still permitting each <code>pty</code> to have a <code>utmpx</code> record (as most applications expect.). The <code>NONROOT_USER</code> macro defines the value that <code>login</code> places in the <code>ut_exit.e_exit</code> member.</p>
RETURN VALUES	<p>Upon successful completion, <code>getutxent ()</code>, <code>getutxid ()</code>, and <code>getutxline ()</code> each return a pointer to a <code>utmpx</code> structure containing a copy of the requested entry in the user accounting database. Otherwise a null pointer is returned.</p> <p>The return value may point to a static area which is overwritten by a subsequent call to <code>getutxid ()</code> or <code>getutxline ()</code>.</p> <p>Upon successful completion, <code>pututxline ()</code> returns a pointer to a <code>utmpx</code> structure containing a copy of the entry added to the user accounting database. Otherwise a null pointer is returned.</p> <p>The <code>endutxent ()</code> and <code>setutxent ()</code> functions return no value.</p> <p>A null pointer is returned upon failure to read, whether for permissions or having reached the end of file, or upon failure to write.</p>
USAGE	<p>These functions use buffered standard I/O for input, but <code>pututxline ()</code> uses an unbuffered write to avoid race conditions between processes trying to modify the <code>utmpx</code> and <code>wtmpx</code> files.</p>

getutxent(3C)

Applications should not access the `utmpx` and `wtmpx` databases directly, but should use these functions to ensure that these databases are maintained consistently.

FILES `/var/adm/utmpx` user access and accounting information
`/var/adm/wtmpx` history of user access and accounting information

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	<code>endutxent()</code> , <code>getutxent()</code> , <code>getutxid()</code> , <code>getutxline()</code> , <code>pututxline()</code> , and <code>setutxent()</code> are Standard.
MT-Level	Unsafe

SEE ALSO `getutent(3C)`, `ttyslot(3C)`, `wait(3C)`, `wait.h(3HEAD)`, `utmpx(4)`, `attributes(5)`, `standards(5)`

NOTES The most current entry is saved in a static structure. Multiple accesses require that it be copied before further accesses are made. On each call to either `getutxid()` or `getutxline()`, the routine examines the static structure before performing more I/O. If the contents of the static structure match what it is searching for, it looks no further. For this reason, to use `getutxline()` to search for multiple occurrences it would be necessary to zero out the static after each success, or `getutxline()` would just return the same structure over and over again. There is one exception to the rule about emptying the structure before further reads are done. The implicit read done by `pututxline()` (if it finds that it is not already at the correct place in the file) will not hurt the contents of the static structure returned by the `getutxent()`, `getutxid()`, or `getutxline()` routines, if the user has just modified those contents and passed the pointer back to `pututxline()`.

NAME	getvfsent, getvfsfile, getvfsspec, getvfsany – get vfstab file entry						
SYNOPSIS	<pre>#include <stdio.h> #include <sys/vfstab.h> int getvfsent(FILE *fp, struct vfstab *vp); int getvfsfile(FILE *fp, struct vfstab *vp, char *file); int getvfsspec(FILE *, struct vfstab *vp, char *spec); int getvfsany(FILE *, struct vfstab *vp, struct vfstab *vref);</pre>						
DESCRIPTION	<p>The <code>getvfsent()</code>, <code>getvfsfile()</code>, <code>getvfsspec()</code>, and <code>getvfsany()</code> functions each fill in the structure pointed to by <code>vp</code> with the broken-out fields of a line in the <code>/etc/vfstab</code> file. Each line in the file contains a <code>vfstab</code> structure, declared in the <code><sys/vfstab.h></code> header, whose following members are described on the <code>vfstab(4)</code> manual page:</p> <pre>char *vfs_special; char *vfs_fsckdev; char *vfs_mountp; char *vfs_fstype; char *vfs_fsckpass; char *vfs_automnt; char *vfs_mntopts;</pre> <p>The <code>getvfsent()</code> function returns a pointer to the next <code>vfstab</code> structure in the file; so successive calls can be used to search the entire file.</p> <p>The <code>getvfsfile()</code> function searches the file referenced by <code>fp</code> until a mount point matching file is found and fills <code>vp</code> with the fields from the line in the file.</p> <p>The <code>getvfsspec()</code> function searches the file referenced by <code>fp</code> until a special device matching <code>spec</code> is found and fills <code>vp</code> with the fields from the line in the file. The <code>spec</code> argument will try to match on device type (block or character special) and major and minor device numbers. If it cannot match in this manner, then it compares the strings.</p> <p>The <code>getvfsany()</code> function searches the file referenced by <code>fp</code> until a match is found between a line in the file and <code>vref</code>. A match occurs if all non-null entries in <code>vref</code> match the corresponding fields in the file.</p> <p>Note that these functions do not open, close, or rewind the file.</p>						
RETURN VALUES	<p>If the next entry is successfully read by <code>getvfsent()</code> or a match is found with <code>getvfsfile()</code>, <code>getvfsspec()</code>, or <code>getvfsany()</code>, 0 is returned. If an end-of-file is encountered on reading, these functions return -1. If an error is encountered, a value greater than 0 is returned. The possible error values are:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>VFS_TOOLONG</code></td> <td>A line in the file exceeded the internal buffer size of <code>VFS_LINE_MAX</code>.</td> </tr> <tr> <td><code>VFS_TOOMANY</code></td> <td>A line in the file contains too many fields.</td> </tr> <tr> <td><code>VFS_TOOFEW</code></td> <td>A line in the file contains too few fields.</td> </tr> </table>	<code>VFS_TOOLONG</code>	A line in the file exceeded the internal buffer size of <code>VFS_LINE_MAX</code> .	<code>VFS_TOOMANY</code>	A line in the file contains too many fields.	<code>VFS_TOOFEW</code>	A line in the file contains too few fields.
<code>VFS_TOOLONG</code>	A line in the file exceeded the internal buffer size of <code>VFS_LINE_MAX</code> .						
<code>VFS_TOOMANY</code>	A line in the file contains too many fields.						
<code>VFS_TOOFEW</code>	A line in the file contains too few fields.						

getvfsent(3C)

FILES | /etc/vfstab

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `vfstab(4)`, `attributes(5)`

NOTES The members of the `vfstab` structure point to information contained in a static area, so it must be copied if it is to be saved.

NAME	getwc – get wide character from a stream						
SYNOPSIS	<pre>#include <stdio.h> #include <wchar.h> wint_t getwc(FILE *<i>stream</i>) ;</pre>						
DESCRIPTION	The <code>getwc()</code> function is equivalent to fgetwc(3C) , except that if it is implemented as a macro it may evaluate <i>stream</i> more than once, so the argument should never be an expression with side effects.						
RETURN VALUES	Refer to fgetwc(3C) .						
ERRORS	Refer to fgetwc(3C) .						
USAGE	<p>This interface is provided to align with some current implementations and with possible future ISO standards.</p> <p>Because it may be implemented as a macro, <code>getwc()</code> may treat incorrectly a <i>stream</i> argument with side effects. In particular, <code>getwc(*f++)</code> may not work as expected. Therefore, use of this function is not recommended; fgetwc(3C) should be used instead.</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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	fgetwc(3C) , attributes(5) , standards(5)						

getwchar(3C)

NAME	getwchar – get wide character from stdin stream						
SYNOPSIS	<pre>#include <wchar.h> wint_t getwchar(void);</pre>						
DESCRIPTION	The <code>getwchar()</code> function is equivalent to <code>getwc(stdin)</code> .						
RETURN VALUES	Refer to fgetwc(3C) .						
ERRORS	Refer to fgetwc(3C) .						
USAGE	If the <code>wint_t</code> value returned by <code>getwchar()</code> is stored into a variable of type <code>wchar_t</code> and then compared against the <code>wint_t</code> macro <code>WEOF</code> , the comparison may never succeed because <code>wchar_t</code> is defined as unsigned.						
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>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	fgetwc(3C) , getwc(3C) , attributes(5) , standards(5)						

NAME	getwd – get current working directory pathname				
SYNOPSIS	<pre>#include <unistd.h> char *getwd(char *path_name);</pre>				
DESCRIPTION	<p>The <code>getwd()</code> function determines an absolute pathname of the current working directory of the calling process, and copies that pathname into the array pointed to by the <code>path_name</code> argument.</p> <p>If the length of the pathname of the current working directory is greater than <code>(PATH_MAX + 1)</code> including the null byte, <code>getwd()</code> fails and returns a null pointer.</p>				
RETURN VALUES	Upon successful completion, a pointer to the string containing the absolute pathname of the current working directory is returned. Otherwise, <code>getwd()</code> returns a null pointer and the contents of the array pointed to by <code>path_name</code> are undefined.				
ERRORS	No errors are defined.				
USAGE	For portability to implementations conforming to versions of the X/Open Portability Guide prior to SUS, <code>getcwd(3C)</code> is preferred over this function.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Standard				
SEE ALSO	<code>getcwd(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>				

getwidth(3C)

NAME	getwidth – get codeset information				
SYNOPSIS	<pre>#include <euc.h> #include <getwidth.h> void getwidth(eucwidth_t *ptr) ;</pre>				
DESCRIPTION	<p>The <code>getwidth()</code> function reads the character class table for the current locale to get information on the supplementary codesets. <code>getwidth()</code> sets this information into the struct <code>eucwidth_t</code>. This struct is defined in <code><euc.h></code> and has the following members:</p> <pre>short int _eucw1, _eucw2, _eucw3; short int _scrw1, _scrw2, _scrw3; short int _pcw; char _multibyte;</pre> <p>Codeset width values for supplementary codesets 1, 2, and 3 are set in <code>_eucw1</code>, <code>_eucw2</code>, and <code>_eucw3</code>, respectively. Screen width values for supplementary codesets 1, 2, and 3 are set in <code>_scrw1</code>, <code>_scrw2</code>, and <code>_scrw3</code>, respectively.</p> <p>The width of Extended Unix Code (EUC) Process Code is set in <code>_pcw</code>. The <code>_multibyte</code> entry is set to 1 if multibyte characters are used, and set to 0 if only single-byte characters are used.</p>				
ATTRIBUTES	See <code>attributes(5)</code> 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>MT-Safe with exceptions</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe with exceptions				
SEE ALSO	<code>euclen(3C)</code> , <code>setlocale(3C)</code> , <code>attributes(5)</code>				
NOTES	<p>The <code>getwidth()</code> function can be used safely in a multithreaded application, as long as <code>setlocale(3C)</code> is not being called to change the locale.</p> <p>The <code>getwidth()</code> function will only work with EUC locales.</p>				

NAME	getws, fgetws – get a wide-character string from a stream
SYNOPSIS	<pre>#include <stdio.h> #include <wchar.h> wchar_t *getws (wchar_t *ws) ; #include <stdio.h> #include <wchar.h> wchar_t *fgetws (wchar_t *restrict ws, int n, FILE *restrict stream) ;</pre>
DESCRIPTION	<p>The <code>getws()</code> function reads a string of characters from the standard input stream, <code>stdin</code>, converts these characters to the corresponding wide-character codes, and writes them to the array pointed to by <code>ws</code>, until a newline character is read, converted and transferred to <code>ws</code> or an end-of-file condition is encountered. The wide-character string, <code>ws</code>, is then terminated with a null wide-character code.</p> <p>The <code>fgetws()</code> function reads characters from the <code>stream</code>, converts them to the corresponding wide-character codes, and places them in the <code>wchar_t</code> array pointed to by <code>ws</code> until <code>n-1</code> characters are read, or until a newline character is read, converted and transferred to <code>ws</code>, or an end-of-file condition is encountered. The wide-character string, <code>ws</code>, is then terminated with a null wide-character code.</p> <p>If an error occurs, the resulting value of the file position indicator for the stream is indeterminate.</p> <p>The <code>fgetws()</code> function may mark the <code>st_atime</code> field of the file associated with <code>stream</code> for update. The <code>st_atime</code> field will be marked for update by the first successful execution of <code>fgetc(3C)</code>, <code>fgets(3C)</code>, <code>fgetwc(3C)</code>, <code>fgetws()</code>, <code>fread(3C)</code>, <code>fscanf(3C)</code>, <code>getc(3C)</code>, <code>getchar(3C)</code>, <code>gets(3C)</code>, or <code>scanf(3C)</code> using <code>stream</code> that returns data not supplied by a prior call to <code>ungetc(3C)</code> or <code>ungetwc(3C)</code>.</p>
RETURN VALUES	<p>Upon successful completion, <code>getws()</code> and <code>fgetws()</code> return <code>ws</code>. If the stream is at end-of-file, the end-of-file indicator for the stream is set and <code>fgetws()</code> returns a null pointer. For standard-conforming (see <code>standards(5)</code>) applications, if the end-of-file indicator for the stream is set, <code>fgetws()</code> returns a null pointer whether or not the stream is at end-of-file. If a read error occurs, the error indicator for the stream is set and <code>fgetws()</code> returns a null pointer and sets <code>errno</code> to indicate the error.</p>
ERRORS	See <code>fgetwc(3C)</code> for the conditions that will cause <code>fgetws()</code> to fail.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	<code>fgetws()</code> is Standard.
MT-Level	MT-Safe

getws(3C)

SEE ALSO [ferror\(3C\)](#), [fgetwc\(3C\)](#), [fread\(3C\)](#), [getwc\(3C\)](#), [putws\(3C\)](#), [scanf\(3C\)](#), [ungetc\(3C\)](#), [ungetwc\(3C\)](#)[attributes\(5\)](#), [standards\(5\)](#)

NAME	getzoneid, getzoneidbyname, getzonenamebyid – map between zone id and name						
SYNOPSIS	<pre>#include <zone.h> zoneid_t getzoneid(void); zoneid_t getzoneidbyname(const char *name); ssize_t getzonenamebyid(zoneid_t id, char *buf, size_t buflen);</pre>						
DESCRIPTION	<p>The <code>getzoneid()</code> function returns the zone ID of the calling process.</p> <p>The <code>getzoneidbyname()</code> function returns the zone ID corresponding to the named zone, if that zone is currently active. If <i>name</i> is <code>NULL</code>, the function returns the zone ID of the calling process.</p> <p>The <code>getzonenamebyid()</code> function stores the name of the zone with ID specified by <i>id</i> in the location specified by <i>buf</i>. The <i>bufsize</i> argument specifies the size in bytes of the buffer. If the buffer is too small to hold the complete null-terminated name, the first <i>bufsize</i> bytes of the name are stored in the buffer. A buffer of size <code>{ZONENAME_MAX}</code> is sufficient to hold any zone name. If <i>buf</i> is <code>NULL</code> or <i>bufsize</i> is 0, the name is not copied into the buffer.</p>						
RETURN VALUES	<p>On successful completion, <code>getzoneid()</code> and <code>getzoneidbyname()</code> return a non-negative zone ID. Otherwise, <code>getzoneidbyname()</code> returns <code>-1</code> and sets <code>errno</code> to indicate the error.</p> <p>On successful completion, the <code>getzonenamebyid()</code> function returns the buffer size required to hold the full null-terminated name. Otherwise, it returns <code>-1</code> and sets <code>errno</code> to indicate the error.</p>						
ERRORS	<p>The <code>getzoneidbyname()</code> function will fail if:</p> <p><code>EFAULT</code> The <i>name</i> argument is non-null and points to an illegal address.</p> <p><code>EINVAL</code> A zone with the indicated <i>name</i> is not active.</p> <p><code>ENAMETOOLONG</code> The length of the <i>name</i> argument exceeds <code>{ZONENAME_MAX}</code>.</p> <p>The <code>getzonenamebyid()</code> function will fail if:</p> <p><code>EINVAL</code> A zone with the specified ID is not active.</p> <p><code>EFAULT</code> The <i>buf</i> argument points to an illegal address.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
MT-Level	Safe						
SEE ALSO	<code>intro(2)</code> , <code>chroot(2)</code> , <code>malloc(3C)</code> , <code>attributes(5)</code> , <code>zones(5)</code>						

glob(3C)

NAME	glob, globfree – generate path names matching a pattern
SYNOPSIS	<pre>#include <glob.h> int glob(const char *restrict pattern, int flags, int (*errfunc)(const char *epath, int eerrno), glob_t *restrict pglob); void globfree(glob_t *pglob);</pre>
DESCRIPTION	<p>The glob() function is a path name generator.</p> <p>The globfree() function frees any memory allocated by glob() associated with pglob.</p>
<i>pattern</i> Argument	<p>The argument <i>pattern</i> is a pointer to a path name pattern to be expanded. The glob() function matches all accessible path names against this pattern and develops a list of all path names that match. In order to have access to a path name, glob() requires search permission on every component of a path except the last, and read permission on each directory of any filename component of <i>pattern</i> that contains any of the following special characters:</p> <p>* ? [</p>
<i>pglob</i> Argument	<p>The structure type glob_t is defined in the header <glob.h> and includes at least the following members:</p> <pre>size_t gl_pathc; /* count of paths matched by pattern */ char **gl_pathv; /* pointer to list of matched path names */ size_t gl_offs; /* slots to reserve at beginning of gl_pathv */</pre> <p>The glob() function stores the number of matched path names into <i>pglob</i>→gl_pathc and a pointer to a list of pointers to path names into <i>pglob</i>→gl_pathv. The path names are in sort order as defined by the current setting of the LC_COLLATE category. The first pointer after the last path name is a NULL pointer. If the pattern does not match any path names, the returned number of matched paths is set to 0, and the contents of <i>pglob</i>→gl_pathv are implementation-dependent.</p> <p>It is the caller's responsibility to create the structure pointed to by <i>pglob</i>. The glob() function allocates other space as needed, including the memory pointed to by gl_pathv. The globfree() function frees any space associated with <i>pglob</i> from a previous call to glob().</p>
<i>flags</i> Argument	<p>The <i>flags</i> argument is used to control the behavior of glob(). The value of <i>flags</i> is a bitwise inclusive OR of zero or more of the following constants, which are defined in the header <glob.h>:</p> <pre> GLOBAL_APPEND Append path names generated to the ones from a previous call to glob(). GLOBAL_DOOFFS Make use of <i>pglob</i>→gl_offs. If this flag is set, <i>pglob</i>→gl_offs is used to specify how many NULL pointers to add to the beginning of <i>pglob</i>→gl_pathv. In other words, <i>pglob</i>→gl_pathv will point</pre>

to *pglob*→*gl_offs* NULL pointers, followed by *pglob*→*gl_pathc* path name pointers, followed by a NULL pointer.

GLOB_ERR	Causes <code>glob()</code> to return when it encounters a directory that it cannot open or read. Ordinarily, <code>glob()</code> continues to find matches.
GLOB_MARK	Each path name that is a directory that matches <i>pattern</i> has a slash appended.
GLOB_NOCHECK	If <i>pattern</i> does not match any path name, then <code>glob()</code> returns a list consisting of only <i>pattern</i> , and the number of matched path names is 1.
GLOB_NOESCAPE	Disable backslash escaping.
GLOB_NOSORT	Ordinarily, <code>glob()</code> sorts the matching path names according to the current setting of the <code>LC_COLLATE</code> category. When this flag is used the order of path names returned is unspecified.

The `GLOB_APPEND` flag can be used to append a new set of path names to those found in a previous call to `glob()`. The following rules apply when two or more calls to `glob()` are made with the same value of *pglob* and without intervening calls to `globfree()`:

1. The first such call must not set `GLOB_APPEND`. All subsequent calls must set it.
2. All the calls must set `GLOB_DOOFFS`, or all must not set it.
3. After the second call, *pglob*→*gl_pathv* points to a list containing the following:
 - a. Zero or more NULL pointers, as specified by `GLOB_DOOFFS` and *pglob*→*gl_offs*.
 - b. Pointers to the path names that were in the *pglob*→*gl_pathv* list before the call, in the same order as before.
 - c. Pointers to the new path names generated by the second call, in the specified order.
4. The count returned in *pglob*→*gl_pathc* will be the total number of path names from the two calls.
5. The application can change any of the fields after a call to `glob()`. If it does, it must reset them to the original value before a subsequent call, using the same *pglob* value, to `globfree()` or `glob()` with the `GLOB_APPEND` flag.

errfunc and *epath* Arguments

If, during the search, a directory is encountered that cannot be opened or read and *errfunc* is not a NULL pointer, `glob()` calls (**errfunc*) with two arguments:

1. The *epath* argument is a pointer to the path that failed.
2. The *eerrno* argument is the value of *errno* from the failure, as set by the `opendir(3C)`, `readdir(3C)` or `stat(2)` functions. (Other values may be used to report other errors not explicitly documented for those functions.)

glob(3C)

The following constants are defined as error return values for `glob()`:

<code>GLOB_ABORTED</code>	The scan was stopped because <code>GLOB_ERR</code> was set or <code>(*errfunc)</code> returned non-zero.
<code>GLOB_NOMATCH</code>	The pattern does not match any existing path name, and <code>GLOB_NOCHECK</code> was not set in flags.
<code>GLOB_NOSPACE</code>	An attempt to allocate memory failed.

If `(*errfunc)` is called and returns non-zero, or if the `GLOB_ERR` flag is set in *flags*, `glob()` stops the scan and returns `GLOB_ABORTED` after setting `gl_pathc` and `gl_pathv` in *pglob* to reflect the paths already scanned. If `GLOB_ERR` is not set and either `errfunc` is a NULL pointer or `(*errfunc)` returns 0, the error is ignored.

RETURN VALUES

The following values are returned by `glob()`:

0	Successful completion. The argument <code>pglob->gl_pathc</code> returns the number of matched path names and the argument <code>pglob->gl_pathv</code> contains a pointer to a null-terminated list of matched and sorted path names. However, if <code>pglob->gl_pathc</code> is 0, the content of <code>pglob->gl_pathv</code> is undefined.
non-zero	An error has occurred. Non-zero constants are defined in <code><glob.h></code> . The arguments <code>pglob->gl_pathc</code> and <code>pglob->gl_pathv</code> are still set as defined above.

The `globfree()` function returns no value.

USAGE

This function is not provided for the purpose of enabling utilities to perform path name expansion on their arguments, as this operation is performed by the shell, and utilities are explicitly not expected to redo this. Instead, it is provided for applications that need to do path name expansion on strings obtained from other sources, such as a pattern typed by a user or read from a file.

If a utility needs to see if a path name matches a given pattern, it can use [fnmatch\(3C\)](#).

Note that `gl_pathc` and `gl_pathv` have meaning even if `glob()` fails. This allows `glob()` to report partial results in the event of an error. However, if `gl_pathc` is 0, `gl_pathv` is unspecified even if `glob()` did not return an error.

The `GLOB_NOCHECK` option could be used when an application wants to expand a path name if wildcards are specified, but wants to treat the pattern as just a string otherwise.

The new path names generated by a subsequent call with `GLOB_APPEND` are not sorted together with the previous path names. This mirrors the way that the shell handles path name expansion when multiple expansions are done on a command line.

Applications that need tilde and parameter expansion should use the [wordexp\(3C\)](#) function.

EXAMPLES **EXAMPLE 1** Example of `glob_dooofs` function.

One use of the `GLOB_DOOFFS` flag is by applications that build an argument list for use with the `execv()`, `execve()`, or `execvp()` functions (see [exec\(2\)](#)). Suppose, for example, that an application wants to do the equivalent of:

```
ls -l *.c
```

but for some reason:

```
system("ls -l *.c")
```

is not acceptable. The application could obtain approximately the same result using the sequence:

```
globbuf.gl_offs = 2;
glob (*.c", GLOB_DOOFFS, NULL, &globbuf);
globbuf.gl_pathv[0] = "ls";
globbuf.gl_pathv[1] = "-l";
execvp ("ls", &globbuf.gl_pathv[0]);
```

Using the same example:

```
ls -l *.c *.h
```

could be approximately simulated using `GLOB_APPEND` as follows:

```
globbuf.gl_offs = 2;
glob (*.c", GLOB_DOOFFS, NULL, &globbuf);
glob (*.h", GLOB_DOOFFS|GLOB_APPEND, NULL, &globbuf);
. . .
```

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [execv\(2\)](#), [stat\(2\)](#), [fnmatch\(3C\)](#), [opendir\(3C\)](#), [readdir\(3C\)](#), [wordexp\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

grantpt(3C)

NAME	grantpt – grant access to the slave pseudo-terminal device						
SYNOPSIS	<pre>#include <stdlib.h> int grantpt(int <i>fildev</i>);</pre>						
DESCRIPTION	The <code>grantpt()</code> function changes the mode and ownership of the slave pseudo-terminal device associated with its master pseudo-terminal counterpart. <i>fildev</i> is the file descriptor returned from a successful open of the master pseudo-terminal device. A <i>setuid</i> root program (see <code>setuid(2)</code>) is invoked to change the permissions. The user ID of the slave is set to the real UID of the calling process and the group ID is set to a reserved group. The permission mode of the slave pseudo-terminal is set to readable and writable by the owner and writable by the group.						
RETURN VALUES	Upon successful completion, <code>grantpt()</code> returns 0. Otherwise, it returns -1 and sets <code>errno</code> to indicate the error.						
ERRORS	The <code>grantpt()</code> function may fail if: EBADF The <i>fildev</i> argument is not a valid open file descriptor. EINVAL The <i>fildev</i> argument is not associated with a master pseudo-terminal device. EACCES The corresponding slave pseudo-terminal device could not be accessed.						
USAGE	The <code>grantpt()</code> function will fail if it is unable to successfully invoke the <i>setuid</i> root program. It may also fail if the application has installed a signal handler to catch SIGCHLD signals.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes: <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	<code>open(2)</code> , <code>setuid(2)</code> , <code>ptsname(3C)</code> , <code>unlockpt(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code> <i>STREAMS Programming Guide</i>						

NAME	hsearch, hcreate, hdestroy – manage hash search tables
SYNOPSIS	<pre>#include <search.h> ENTRY *hsearch(ENTRY item, ACTION action); int hcreate(size_t mekments); void hdestroy(void);</pre>
DESCRIPTION	<p>The <code>hsearch()</code> function is a hash-table search routine generalized from Knuth (6.4) Algorithm D. It returns a pointer into a hash table indicating the location at which an entry can be found. The comparison function used by <code>hsearch()</code> is <code>strcmp()</code> (see string(3C)). The <code>item</code> argument is a structure of type <code>ENTRY</code> (defined in the <code><search.h></code> header) containing two pointers: <code>item.key</code> points to the comparison key, and <code>item.data</code> points to any other data to be associated with that key. (Pointers to types other than <code>void</code> should be cast to <code>pointer-to-void</code>.) The <code>action</code> argument is a member of an enumeration type <code>ACTION</code> (defined in <code><search.h></code>) indicating the disposition of the entry if it cannot be found in the table. <code>ENTER</code> indicates that the item should be inserted in the table at an appropriate point. Given a duplicate of an existing item, the new item is not entered and <code>hsearch()</code> returns a pointer to the existing item. <code>FIND</code> indicates that no entry should be made. Unsuccessful resolution is indicated by the return of a null pointer.</p> <p>The <code>hcreate()</code> function allocates sufficient space for the table, and must be called before <code>hsearch()</code> is used. The <code>nel</code> argument is an estimate of the maximum number of entries that the table will contain. This number may be adjusted upward by the algorithm in order to obtain certain mathematically favorable circumstances.</p> <p>The <code>hdestroy()</code> function destroys the search table, and may be followed by another call to <code>hcreate()</code>.</p>
RETURN VALUES	<p>The <code>hsearch()</code> function returns a null pointer if either the action is <code>FIND</code> and the item could not be found or the action is <code>ENTER</code> and the table is full.</p> <p>The <code>hcreate()</code> function returns 0 if it cannot allocate sufficient space for the table.</p>
USAGE	<p>The <code>hsearch()</code> and <code>hcreate()</code> functions use malloc(3C) to allocate space.</p> <p>Only one hash search table may be active at any given time.</p>
EXAMPLES	<p>EXAMPLE 1 Example to read in strings.</p> <p>The following example will read in strings followed by two numbers and store them in a hash table, discarding duplicates. It will then read in strings and find the matching entry in the hash table and print it.</p> <pre>#include <stdio.h> #include <search.h> #include <string.h> #include <stdlib.h> struct info { /* this is the info stored in table */</pre>

EXAMPLE 1 Example to read in strings. (Continued)

```

        int age, room;                /* other than the key */
};
#define NUM_EMPL    5000             /* # of elements in search table */
main( )
{
    /* space to store strings */
    char string_space[NUM_EMPL*20];
    /* space to store employee info */
    struct info info_space[NUM_EMPL];
    /* next avail space in string_space */
    char *str_ptr = string_space;
    /* next avail space in info_space */
    struct info *info_ptr = info_space;
    ENTRY item, *found_item;
    /* name to look for in table */
    char name_to_find[30];
    int i = 0;

    /* create table */
    (void) hcreate(NUM_EMPL);
    while (scanf("%s%d%d", str_ptr, &info_ptr->age,
        &info_ptr->room) != EOF && i++ < NUM_EMPL) {
        /* put info in structure, and structure in item */
        item.key = str_ptr;
        item.data = (void *)info_ptr;
        str_ptr += strlen(str_ptr) + 1;
        info_ptr++;
        /* put item into table */
        (void) hsearch(item, ENTER);
    }

    /* access table */
    item.key = name_to_find;
    while (scanf("%s", item.key) != EOF) {
        if ((found_item = hsearch(item, FIND)) != NULL) {
            /* if item is in the table */
            (void)printf("found %s, age = %d, room = %d\n",
                found_item->key,
                ((struct info *)found_item->data)->age,
                ((struct info *)found_item->data)->room);
        } else {
            (void)printf("no such employee %s\n",
                name_to_find)
        }
    }
    return 0;
}

```

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

hsearch(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Safe

SEE ALSO [bsearch\(3C\)](#), [lsearch\(3C\)](#), [malloc\(3C\)](#), [string\(3C\)](#), [tsearch\(3C\)](#), [malloc\(3MALLOC\)](#), [attributes\(5\)](#), [standards\(5\)](#)

The Art of Computer Programming, Volume 3, Sorting and Searching by Donald E. Knuth, published by Addison-Wesley Publishing Company, 1973.

iconv(3C)

NAME	iconv – code conversion function
SYNOPSIS	
Default	<pre>#include <iconv.h> extern size_t iconv(iconv_t <i>cd</i>, const char **restrict <i>inbuf</i>, size_t *restrict <i>inbytesleft</i>, char **restrict <i>outbuf</i>, size_t *restrict <i>outbytesleft</i>);</pre>
SUSv3	<pre>#include <iconv.h> size_t iconv(iconv_t <i>cd</i>, char **restrict <i>inbuf</i>, size_t *restrict <i>inbytesleft</i>, char **restrict <i>outbuf</i>, size_t *restrict <i>outbytesleft</i>);</pre>
DESCRIPTION	<p>The <code>iconv()</code> function converts the sequence of characters from one code set, in the array specified by <i>inbuf</i>, into a sequence of corresponding characters in another code set, in the array specified by <i>outbuf</i>. The code sets are those specified in the <code>iconv_open()</code> call that returned the conversion descriptor, <i>cd</i>. The <i>inbuf</i> argument points to a variable that points to the first character in the input buffer and <i>inbytesleft</i> indicates the number of bytes to the end of the buffer to be converted. The <i>outbuf</i> argument points to a variable that points to the first available byte in the output buffer and <i>outbytesleft</i> indicates the number of the available bytes to the end of the buffer.</p> <p>For state-dependent encodings, the conversion descriptor <i>cd</i> is placed into its initial shift state by a call for which <i>inbuf</i> is a null pointer, or for which <i>inbuf</i> points to a null pointer. When <code>iconv()</code> is called in this way, and if <i>outbuf</i> is not a null pointer or a pointer to a null pointer, and <i>outbytesleft</i> points to a positive value, <code>iconv()</code> will place, into the output buffer, the byte sequence to change the output buffer to its initial shift state. If the output buffer is not large enough to hold the entire reset sequence, <code>iconv()</code> will fail and set <code>errno</code> to <code>E2BIG</code>. Subsequent calls with <i>inbuf</i> as other than a null pointer or a pointer to a null pointer cause the conversion to take place from the current state of the conversion descriptor.</p> <p>If a sequence of input bytes does not form a valid character in the specified code set, conversion stops after the previous successfully converted character. If the input buffer ends with an incomplete character or shift sequence, conversion stops after the previous successfully converted bytes. If the output buffer is not large enough to hold the entire converted input, conversion stops just prior to the input bytes that would cause the output buffer to overflow. The variable pointed to by <i>inbuf</i> is updated to point to the byte following the last byte successfully used in the conversion. The value pointed to by <i>inbytesleft</i> is decremented to reflect the number of bytes still not converted in the input buffer. The variable pointed to by <i>outbuf</i> is updated to point to the byte following the last byte of converted output data. The value pointed to by <i>outbytesleft</i> is decremented to reflect the number of bytes still available in the output buffer. For state-dependent encodings, the conversion descriptor is updated to reflect the shift state in effect at the end of the last successfully converted byte sequence.</p> <p>If <code>iconv()</code> encounters a character in the input buffer that is legal, but for which an identical character does not exist in the target code set, <code>iconv()</code> performs an implementation-defined conversion on this character.</p>

RETURN VALUES The `iconv()` function updates the variables pointed to by the arguments to reflect the extent of the conversion and returns the number of non-identical conversions performed. If the entire string in the input buffer is converted, the value pointed to by *inbytesleft* will be 0. If the input conversion is stopped due to any conditions mentioned above, the value pointed to by *inbytesleft* will be non-zero and `errno` is set to indicate the condition. If an error occurs `iconv()` returns `(size_t) -1` and sets `errno` to indicate the error.

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

EILSEQ	Input conversion stopped due to an input byte that does not belong to the input code set.
E2BIG	Input conversion stopped due to lack of space in the output buffer.
EINVAL	Input conversion stopped due to an incomplete character or shift sequence at the end of the input buffer.

The `iconv()` function may fail if:

EBADF	The <i>cd</i> argument is not a valid open conversion descriptor.
-------	---

EXAMPLES **EXAMPLE 1** Using the `iconv()` Functions

The following example uses the `iconv()` functions:

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

/*
 * For state-dependent encodings, changes the state of the conversion
 * descriptor to initial shift state. Also, outputs the byte sequence
 * to change the state to initial state.
 * This code is assuming the iconv call for initializing the state
 * won't fail due to lack of space in the output buffer.
 */
#define INIT_SHIFT_STATE(cd, fptr, ileft, tptr, oleft) \
{ \
    fptr = NULL; \
    ileft = 0; \
    tptr = to; \
    oleft = BUFSIZ; \
    (void) iconv(cd, &fptr, &ileft, &tptr, &oleft); \
    (void) fwrite(to, 1, BUFSIZ - oleft, stdout); \
}

int
main(int argc, char **argv)
{
    iconv_t cd;
    char    from[BUFSIZ], to[BUFSIZ];
    char    *from_code, *to_code;
    char    *tptr;
```

EXAMPLE 1 Using the iconv() Functions (Continued)

```

const char *fptr;
size_t ileft, oleft, num, ret;

if (argc != 3) {
    (void) fprintf(stderr,
        "Usage: %s from_codeset to_codeset\\n", argv[0]);
    return (1);
}

from_code = argv[1];
to_code = argv[2];

cd = iconv_open((const char *)to_code, (const char *)from_code);
if (cd == (iconv_t)-1) {
    /*
     * iconv_open failed
     */
    (void) fprintf(stderr,
        "iconv_open(%s, %s) failed\\n", to_code, from_code);
    return (1);
}

ileft = 0;
while ((ileft +=
    (num = fread(from + ileft, 1, BUFSIZ - ileft, stdin)) > 0) {
    if (num == 0) {
        /*
         * Input buffer still contains incomplete character
         * or sequence. However, no more input character.
         */

        /*
         * Initializes the conversion descriptor and outputs
         * the sequence to change the state to initial state.
         */
        INIT_SHIFT_STATE(cd, fptr, ileft, tptr, oleft);
        (void) iconv_close(cd);

        (void) fprintf(stderr, "Conversion error\\n");
        return (1);
    }

    fptr = from;
    for (;;) {
        tptr = to;
        oleft = BUFSIZ;

        ret = iconv(cd, &fptr, &ileft, &tptr, &oleft);
        if (ret != (size_t)-1) {
            /*
             * iconv succeeded
             */

```

EXAMPLE 1 Using the iconv() Functions (Continued)

```

        /*
        * Outputs converted characters
        */
        (void) fwrite(to, 1, BUFSIZ - oleft, stdout);
        break;
    }

    /*
    * iconv failed
    */
    if (errno == EINVAL) {
        /*
        * Incomplete character or shift sequence
        */

        /*
        * Outputs converted characters
        */
        (void) fwrite(to, 1, BUFSIZ - oleft, stdout);
        /*
        * Copies remaining characters in input buffer
        * to the top of the input buffer.
        */
        (void) memmove(from, fptr, ileft);
        /*
        * Tries to fill input buffer from stdin
        */
        break;
    } else if (errno == E2BIG) {
        /*
        * Lack of space in output buffer
        */

        /*
        * Outputs converted characters
        */
        (void) fwrite(to, 1, BUFSIZ - oleft, stdout);
        /*
        * Tries to convert remaining characters in
        * input buffer with emptied output buffer
        */
        continue;
    } else if (errno == EILSEQ) {
        /*
        * Illegal character or shift sequence
        */

        /*
        * Outputs converted characters
        */
        (void) fwrite(to, 1, BUFSIZ - oleft, stdout);
        /*
        * Initializes the conversion descriptor and
        * outputs the sequence to change the state to

```

iconv(3C)

EXAMPLE 1 Using the iconv() Functions (Continued)

```
        * initial state.
        */
        INIT_SHIFT_STATE(cd, fptr, ileft, tptr, oleft);
        (void) iconv_close(cd);

        (void) fprintf(stderr,
            "Illegal character or sequence\\n");
        return (1);
    } else if (errno == EBADF) {
        /*
         * Invalid conversion descriptor.
         * Actually, this shouldn't happen here.
         */
        (void) fprintf(stderr, "Conversion error\\n");
        return (1);
    } else {
        /*
         * This errno is not defined
         */
        (void) fprintf(stderr, "iconv error\\n");
        return (1);
    }
}

/*
 * Initializes the conversion descriptor and outputs
 * the sequence to change the state to initial state.
 */
INIT_SHIFT_STATE(cd, fptr, ileft, tptr, oleft);

(void) iconv_close(cd);
return (0);
}
```

FILES

```
/usr/lib/iconv/*.so
conversion modules for 32-bit

/usr/lib/iconv/sparcv9/*.so
conversion modules for 64-bit sparc

/usr/lib/iconv/amd64/*.so
conversion modules for 64-bit amd64

/usr/lib/iconv/geniconvtbl/binarytables/*.bt
conversion binary tables
```

ATTRIBUTES

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [geniconvtbl\(1\)](#), [iconv\(1\)](#), [iconv_close\(3C\)](#), [iconv_open\(3C\)](#), [geniconvtbl\(4\)](#), [attributes\(5\)](#), [iconv\(5\)](#), [iconv_unicode\(5\)](#), [standards\(5\)](#)

iconv_close(3C)

NAME iconv_close – code conversion deallocation function

SYNOPSIS

```
#include <iconv.h>

int iconv_close(iconv_t cd);
```

DESCRIPTION The iconv_close() function deallocates the conversion descriptor cd and all other associated resources allocated by the iconv_open(3C) function.

If a file descriptor is used to implement the type iconv_t, that file descriptor will be closed.

For examples using the iconv_close() function, see iconv(3C).

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

ERRORS The iconv_close() function may fail if:

EBADF The conversion descriptor is invalid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO iconv(3C), iconv_open(3C), attributes(5), standards(5)

NAME	iconv_open – code conversion allocation function								
SYNOPSIS	<pre>#include <iconv.h> iconv_t iconv_open(const char *tocode, const char *fromcode);</pre>								
DESCRIPTION	<p>The <code>iconv_open()</code> function returns a conversion descriptor that describes a conversion from the codeset specified by the string pointed to by the <i>fromcode</i> argument to the codeset specified by the string pointed to by the <i>tocode</i> argument. For state-dependent encodings, the conversion descriptor will be in a codeset-dependent initial shift state, ready for immediate use with the <code>iconv(3C)</code> function.</p> <p>Settings of <i>fromcode</i> and <i>tocode</i> and their permitted combinations are implementation-dependent.</p> <p>The <code>iconv_open()</code> function supports the alias of the encoding name specified in <i>tocode</i> and <i>fromcode</i>. The alias table of the encoding name is described in the file <code>/usr/lib/iconv/alias</code>. See <code>alias(4)</code>.</p> <p>A conversion descriptor remains valid in a process until that process closes it.</p> <p>For examples using the <code>iconv_open()</code> function, see <code>iconv(3C)</code>.</p>								
RETURN VALUES	Upon successful completion <code>iconv_open()</code> returns a conversion descriptor for use on subsequent calls to <code>iconv()</code> . Otherwise, <code>iconv_open()</code> returns (<code>iconv_t</code>) <code>-1</code> and sets <code>errno</code> to indicate the error.								
ERRORS	<p>The <code>iconv_open</code> function may fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EMFILE</td> <td>{ OPEN_MAX } files descriptors are currently open in the calling process.</td> </tr> <tr> <td>ENFILE</td> <td>Too many files are currently open in the system.</td> </tr> <tr> <td>ENOMEM</td> <td>Insufficient storage space is available.</td> </tr> <tr> <td>EINVAL</td> <td>The conversion specified by <i>fromcode</i> and <i>tocode</i> is not supported by the implementation.</td> </tr> </table>	EMFILE	{ OPEN_MAX } files descriptors are currently open in the calling process.	ENFILE	Too many files are currently open in the system.	ENOMEM	Insufficient storage space is available.	EINVAL	The conversion specified by <i>fromcode</i> and <i>tocode</i> is not supported by the implementation.
EMFILE	{ OPEN_MAX } files descriptors are currently open in the calling process.								
ENFILE	Too many files are currently open in the system.								
ENOMEM	Insufficient storage space is available.								
EINVAL	The conversion specified by <i>fromcode</i> and <i>tocode</i> is not supported by the implementation.								
FILES	<code>/usr/lib/iconv/alias</code> alias table file of the encoding name								
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:								
	<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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe		
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Interface Stability	Standard								
MT-Level	MT-Safe								
SEE ALSO	<code>exec(2)</code> , <code>iconv(3C)</code> , <code>iconv_close(3C)</code> , <code>malloc(3C)</code> , <code>alias(4)</code> , <code>attributes(5)</code>								

iconv_open(3C)

NOTES | The `iconv_open()` function uses `malloc(3C)` to allocate space for internal buffer areas. `iconv_open()` may fail if there is insufficient storage space to accommodate these buffers.

Portable applications must assume that conversion descriptors are not valid after a call to one of the `exec` functions (see `exec(2)`).

NAME imaxabs – return absolute value

SYNOPSIS `#include <inttypes.h>`
`intmax_t imaxabs(intmax_t j);`

DESCRIPTION The `imaxabs()` function computes the absolute value of an integer *j*. If the result cannot be represented, the behavior is undefined.

RETURN VALUES The `imaxabs()` function returns the absolute value.

ERRORS No errors are defined.

USAGE The absolute value of the most negative number cannot be represented in two's complement.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [imaxdiv\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

imaxdiv(3C)

NAME	imaxdiv – return quotient and remainder						
SYNOPSIS	<pre>#include <inttypes.h> imaxdiv_t imaxdiv(imaxdiv_t <i>numer</i>, imaxdiv_t <i>denom</i>);</pre>						
DESCRIPTION	The <code>imaxdiv()</code> function computes <i>numer</i> / <i>denom</i> and <i>numer</i> % <i>denom</i> in a single operation.						
RETURN VALUES	The <code>imaxdiv()</code> function returns a structure of type <code>imaxdiv_t</code> , comprising both the quotient and the remainder. The structure contains (in either order) the members <code>quot</code> (the quotient) and <code>rem</code> (the remainder), each of which has type <code>intmax_t</code> . If either part of the result cannot be represented, the behavior is undefined.						
ERRORS	No errors are defined.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>imaxabs(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	index, rindex – string operations				
SYNOPSIS	<pre>#include <strings.h> char *index(const char *s, int c); char *rindex(const char *s, int c);</pre>				
DESCRIPTION	<p>The <code>index()</code> and <code>rindex()</code> functions operate on null-terminated strings.</p> <p>The <code>index()</code> function returns a pointer to the first occurrence of character <code>c</code> in string <code>s</code>.</p> <p>The <code>rindex()</code> function returns a pointer to the last occurrence of character <code>c</code> in string <code>s</code>.</p> <p>Both <code>index()</code> and <code>rindex()</code> return a null pointer if <code>c</code> does not occur in the string. The null character terminating a string is considered to be part of the string.</p>				
USAGE	<p>On most modern computer systems, you can <i>not</i> use a null pointer to indicate a null string. A null pointer is an error and results in an abort of the program. If you wish to indicate a null string, you must use a pointer that points to an explicit null string. On some machines and with some implementations of the C programming language, a null pointer, if dereferenced, would yield a null string. Though often used, this practice is not always portable. Programmers using a null pointer to represent an empty string should be aware of this portability issue. Even on machines where dereferencing a null pointer does not cause an abort of the program, it does not necessarily yield a null string.</p>				
ATTRIBUTES	<p>See <code>attributes(5)</code> 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>Interface Stability</td> <td>Standard</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Standard				
SEE ALSO	<code>bstring(3C)</code> , <code>malloc(3C)</code> , <code>string(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>				

initgroups(3C)

NAME	initgroups – initialize the supplementary group access list				
SYNOPSIS	<pre>#include <grp.h> #include <sys/types.h> int initgroups(const char *name, gid_t basegid);</pre>				
DESCRIPTION	<p>The <code>initgroups()</code> function reads the group database to get the group membership for the user specified by <i>name</i>, and initializes the supplementary group access list of the calling process (see getgrnam(3C) and getgroups(2)). The <i>basegid</i> group ID is also included in the supplementary group access list. This is typically the real group ID from the user database.</p> <p>While scanning the group database, if the number of groups, including the <i>basegid</i> entry, exceeds <code>NGROUPS_MAX</code>, subsequent group entries are ignored.</p>				
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.				
ERRORS	The <code>initgroups()</code> function will fail and not change the supplementary group access list if:				
	<table><tr><td><code>EPERM</code></td><td>The <code>{PRIV_PROC_SETID}</code> privilege is not asserted in the effective set of the calling process.</td></tr></table>	<code>EPERM</code>	The <code>{PRIV_PROC_SETID}</code> privilege is not asserted in the effective set of the calling process.		
<code>EPERM</code>	The <code>{PRIV_PROC_SETID}</code> privilege is not asserted in the effective set of the calling process.				
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>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	getgroups(2) , getgrnam(3C) , attributes(5)				

NAME insque, remque – insert/remove element from a queue

SYNOPSIS

```
include <search.h>
void insque(struct qelem *elem, struct qelem *pred);
void remque(struct qelem *elem);
```

DESCRIPTION The `insque()` and `remque()` functions manipulate queues built from doubly linked lists. Each element in the queue must be in the following form:

```
struct qelem {
    struct qelem *q_forw;
    struct qelem *q_back;
    char q_data[ ];
};
```

The `insque()` function inserts *elem* in a queue immediately after *pred*. The `remque()` function removes an entry *elem* from a queue.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Unsafe

SEE ALSO `attributes(5)`, `standards(5)`

isaexec(3C)

NAME	isaexec – invoke isa-specific executable						
SYNOPSIS	<pre>#include <unistd.h> int isaexec(const char *path, char *const argv[], char *const envp[]);</pre>						
DESCRIPTION	The <code>isaexec()</code> function takes the path specified as <i>path</i> and breaks it into directory and file name components. It enquires from the running system the list of supported instruction set architectures; see <code>isalist(5)</code> . The function traverses the list for an executable file in named subdirectories of the original directory. When such a file is located, <code>execve()</code> is invoked with <i>argv[]</i> and <i>envp[]</i> . See <code>exec(2)</code> .						
RETURN VALUES	If no file is located, <code>isaexec()</code> returns <code>ENOENT</code> . Other return values are the same as for <code>execve()</code> .						
EXAMPLES	<p>EXAMPLE 1 Example of <code>isaexec()</code> function.</p> <p>On a system whose <code>isalist</code> is</p> <pre>sparcv7 sparc</pre> <p>the program</p> <pre>int main(int argc, char *argv[], char *envp[]) { return (isaexec("/bin/thing", argv, envp)); }</pre> <p>will look first for an executable file named <code>/bin/sparcv7/thing</code>, then for an executable file named <code>/bin/sparc/thing</code>. It will invoke <code>execve()</code> on the first executable file it finds named <code>thing</code>.</p> <p>On that same system, a program called <code>/u/bin/tofu</code> can cause either <code>/u/bin/sparcv7/tofu</code> or <code>/u/bin/sparc/tofu</code> to be invoked using the following code:</p> <pre>int main(int argc, char *argv[], char *envp[]) { return (isaexec(getexecname(), argv, envp)); }</pre>						
ATTRIBUTES	See <code>attributes(5)</code> 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>Safe</td> </tr> <tr> <td>Interface Stability</td> <td>Stable</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe	Interface Stability	Stable
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
MT-Level	Safe						
Interface Stability	Stable						
SEE ALSO	<code>exec(2)</code> , <code>getexecname(3C)</code> , <code>attributes(5)</code> , <code>isalist(5)</code>						

NAME	isastream – test a file descriptor
SYNOPSIS	<pre>#include <stropts.h> int isastream(int <i>fildev</i>);</pre>
DESCRIPTION	The <code>isastream()</code> function determines if a file descriptor represents a STREAMS file. The <i>fildev</i> argument refers to an open file descriptor.
RETURN VALUES	Upon successful completion, <code>isastream()</code> returns 1 if <i>fildev</i> represents a STREAMS file, and 0 if it does not. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.
ERRORS	The <code>isastream()</code> function will fail if: EBADF The <i>fildev</i> argument is not a valid file descriptor.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `attributes(5)`, `standards(5)`, `streamio(7I)`
STREAMS Programming Guide

isatty(3C)

NAME	isatty – test for a terminal device						
SYNOPSIS	<pre>#include <unistd.h> int isatty(int <i>fildev</i>);</pre>						
DESCRIPTION	The <code>isatty()</code> function tests whether <i>fildev</i> , an open file descriptor, is associated with a terminal device.						
RETURN VALUES	The <code>isatty()</code> function returns 1 if <i>fildev</i> is associated with a terminal; otherwise it returns 0 and may set <code>errno</code> to indicate the error.						
ERRORS	The <code>isatty()</code> function may fail if: EBADF The <i>fildev</i> argument is not a valid open file descriptor. ENOTTY The <i>fildev</i> argument is not associated with a terminal.						
USAGE	The <code>isatty()</code> function does not necessarily indicate that a human being is available for interaction via <i>fildev</i> . It is quite possible that non-terminal devices are connected to the communications line.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes: <table border="1" data-bbox="444 974 1414 1108"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>ttyname(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	isnand, isnanf, finite, fpclass, unordered – determine type of floating-point number																				
SYNOPSIS	<pre>#include <ieeefp.h> int isnand(double <i>dsrc</i>); int isnanf(float <i>fsrc</i>); int finite(double <i>dsrc</i>); fpclass_t fpclass(double <i>dsrc</i>); int unordered(double <i>dsrc1</i>, double <i>dsrc2</i>);</pre>																				
DESCRIPTION	<p>The <code>isnand()</code> and <code>isnanf()</code> functions return <code>TRUE</code> (1) if the argument <i>dsrc</i> or <i>fsrc</i> is a NaN; otherwise they return <code>FALSE</code> (0).</p> <p>The <code>fpclass()</code> function returns one of the following classes to which <i>dsrc</i> belongs:</p> <table border="0"> <tr><td>FP_SNAN</td><td>signaling NaN</td></tr> <tr><td>FP_QNAN</td><td>quiet NaN</td></tr> <tr><td>FP_NINF</td><td>negative infinity</td></tr> <tr><td>FP_PINF</td><td>positive infinity</td></tr> <tr><td>FP_NDENORM</td><td>negative denormalized non-zero</td></tr> <tr><td>FP_PDENORM</td><td>positive denormalized non-zero</td></tr> <tr><td>FP_NZERO</td><td>negative zero</td></tr> <tr><td>FP_PZERO</td><td>positive zero</td></tr> <tr><td>FP_NNORM</td><td>negative normalized non-zero</td></tr> <tr><td>FP_PNORM</td><td>positive normalized non-zero</td></tr> </table> <p>The <code>finite()</code> function returns <code>TRUE</code> (1) if the argument <i>dsrc</i> is neither infinity nor NaN; otherwise it returns <code>FALSE</code> (0).</p> <p>The <code>unordered()</code> function returns <code>TRUE</code> (1) if one of its two arguments is unordered with respect to the other argument. This is equivalent to reporting whether either argument is NaN. If neither argument is NaN, <code>FALSE</code> (0) is returned.</p> <p>None of these functions generates an exception, even for signaling NaNs.</p>	FP_SNAN	signaling NaN	FP_QNAN	quiet NaN	FP_NINF	negative infinity	FP_PINF	positive infinity	FP_NDENORM	negative denormalized non-zero	FP_PDENORM	positive denormalized non-zero	FP_NZERO	negative zero	FP_PZERO	positive zero	FP_NNORM	negative normalized non-zero	FP_PNORM	positive normalized non-zero
FP_SNAN	signaling NaN																				
FP_QNAN	quiet NaN																				
FP_NINF	negative infinity																				
FP_PINF	positive infinity																				
FP_NDENORM	negative denormalized non-zero																				
FP_PDENORM	positive denormalized non-zero																				
FP_NZERO	negative zero																				
FP_PZERO	positive zero																				
FP_NNORM	negative normalized non-zero																				
FP_PNORM	positive normalized non-zero																				
RETURN VALUES	See DESCRIPTION.																				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:																				

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

isnan(3C)

SEE ALSO [fpgetround\(3C\)](#), [isnan\(3M\)](#), [attributes\(5\)](#)

NAME	iswalpha, isenglish, isideogram, isnumber, isphonogram, isspecial, iswalnum, iswascii, iswblank, iswcntrl, iswdigit, iswgraph, iswlower, iswprint, iswpunct, iswspace, iswupper, iswxdigit – wide-character code classification functions				
SYNOPSIS	<pre>#include <wchar.h> #include <wctype.h> int iswalpha(wint_t <i>wc</i>); int isenglish(wint_t <i>wc</i>); int isideogram(wint_t <i>wc</i>); int isnumber(wint_t <i>wc</i>); int isphonogram(wint_t <i>wc</i>); int isspecial(wint_t <i>wc</i>); int iswalnum(wint_t <i>wc</i>); int iswascii(wint_t <i>wc</i>); int iswblank(wint_t <i>wc</i>); int iswcntrl(wint_t <i>wc</i>); int iswdigit(wint_t <i>wc</i>); int iswgraph(wint_t <i>wc</i>); int iswlower(wint_t <i>wc</i>); int iswprint(wint_t <i>wc</i>); int iswpunct(wint_t <i>wc</i>); int iswspace(wint_t <i>wc</i>); int iswupper(wint_t <i>wc</i>); int iswxdigit(wint_t <i>wc</i>);</pre>				
DESCRIPTION	<p>These functions test whether <i>wc</i> is a wide-character code representing a character of a particular class defined in the LC_CTYPE category of the current locale.</p> <p>In all cases, <i>wc</i> is a wint_t, the value of which must be a wide-character code corresponding to a valid character in the current locale or must equal the value of the macro WEOF. If the argument has any other values, the behavior is undefined.</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">iswalpha(<i>wc</i>)</td> <td>Tests whether <i>wc</i> is a wide-character code representing a character of class "alpha" in the program's current locale.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">isenglish(<i>wc</i>)</td> <td>Tests whether <i>wc</i> is a wide-character code representing an English language character, excluding ASCII characters.</td> </tr> </table>	iswalpha(<i>wc</i>)	Tests whether <i>wc</i> is a wide-character code representing a character of class "alpha" in the program's current locale.	isenglish(<i>wc</i>)	Tests whether <i>wc</i> is a wide-character code representing an English language character, excluding ASCII characters.
iswalpha(<i>wc</i>)	Tests whether <i>wc</i> is a wide-character code representing a character of class "alpha" in the program's current locale.				
isenglish(<i>wc</i>)	Tests whether <i>wc</i> is a wide-character code representing an English language character, excluding ASCII characters.				

iswalpha(3C)

<code>isideogram(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing an ideographic language character, excluding ASCII characters.
<code>isnumber(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing digit [0–9], excluding ASCII characters.
<code>isphonogram(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing a phonetic language character, excluding ASCII characters.
<code>isspecial(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing a special language character, excluding ASCII characters.
<code>iswalnum(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing a character of class "alpha" or "digit" in the program's current locale.
<code>iswascii(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing an ASCII character.
<code>iswblank(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing a character of class "blank" in the program's current locale. This function is not available to applications conforming to standards prior to SUSv3. See <code>standards(5)</code> .
<code>iswlower(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing a character of class "lower" in the program's current locale.
<code>iswcntrl(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing a character of class "cntrl" in the program's current locale.
<code>iswdigit(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing a character of class "digit" in the program's current locale.
<code>iswgraph(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing a character of class "graph" in the program's current locale.
<code>iswprint(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing a character of class "print" in the program's current locale.
<code>iswpunct(<i>wc</i>)</code>	Tests whether <i>wc</i> is a wide-character code representing a character of class "punct" in the program's current locale.

iswalpha(3C)

`isspace(wc)` Tests whether *wc* is a wide-character code representing a character of class "space" in the program's current locale.

`iswupper(wc)` Tests whether *wc* is a wide-character code representing a character of class "upper" in the program's current locale.

`iswxdigit(wc)` Tests whether *wc* is a wide-character code representing a character of class "xdigit" in the program's current locale.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	See below.
MT-Level	MT-Safe with exceptions

The `iswalpha()`, `iswalnum()`, `iswblank()`, `iswcntrl()`, `iswdigit()`, `iswgraph()`, `iswlower()`, `iswprint()`, `iswpunct()`, `isspace()`, `iswupper()`, and `iswxdigit()` functions are Standard.

SEE ALSO `localedef(1)`, `setlocale(3C)`, `stdio(3C)`, `ascii(5)`, `attributes(5)`, `standards(5)`

iswctype(3C)

NAME	iswctype – test character for specified class
SYNOPSIS	<pre>#include <wchar.h> int iswctype(wint_t <i>wc</i>, wctype_t <i>charclass</i>);</pre>
DESCRIPTION	The <code>iswctype()</code> function determines whether the wide-character code <i>wc</i> has the character class <i>charclass</i> , returning TRUE or FALSE. The <code>iswctype()</code> function is defined on WEOF and wide-character codes corresponding to the valid character encodings in the current locale. If the <i>wc</i> argument is not in the domain of the function, the result is undefined. If the value of <i>charclass</i> is invalid (that is, not obtained by a call to <code>wctype(3C)</code> or <i>charclass</i> is invalidated by a subsequent call to <code>setlocale(3C)</code> that has affected category LC_CTYPE), the result is indeterminate.
RETURN VALUES	The <code>iswctype()</code> function returns 0 for FALSE and non-zero for TRUE.
USAGE	There are twelve strings that are reserved for the standard character classes:

"alnum"	"alpha"	"blank"
"cntrl"	"digit"	"graph"
"lower"	"print"	"punct"
"space"	"upper"	"xdigit"

In the table below, the functions in the left column are equivalent to the functions in the right column.

<code>iswalnum(<i>wc</i>)</code>	<code>iswctype(<i>wc</i>, wctype("alnum"))</code>
<code>iswalpha(<i>wc</i>)</code>	<code>iswctype(<i>wc</i>, wctype("alpha"))</code>
<code>iswcntrl(<i>wc</i>)</code>	<code>iswctype(<i>wc</i>, wctype("cntrl"))</code>
<code>iswdigit(<i>wc</i>)</code>	<code>iswctype(<i>wc</i>, wctype("digit"))</code>
<code>iswgraph(<i>wc</i>)</code>	<code>iswctype(<i>wc</i>, wctype("graph"))</code>
<code>iswlower(<i>wc</i>)</code>	<code>iswctype(<i>wc</i>, wctype("lower"))</code>
<code>iswprint(<i>wc</i>)</code>	<code>iswctype(<i>wc</i>, wctype("print"))</code>
<code>iswpunct(<i>wc</i>)</code>	<code>iswctype(<i>wc</i>, wctype("punct"))</code>
<code>iswspace(<i>wc</i>)</code>	<code>iswctype(<i>wc</i>, wctype("space"))</code>
<code>iswupper(<i>wc</i>)</code>	<code>iswctype(<i>wc</i>, wctype("upper"))</code>
<code>iswxdigit(<i>wc</i>)</code>	<code>iswctype(<i>wc</i>, wctype("xdigit"))</code>

The call

`iswctype(wc, wctype("blank"))`

does not have an equivalent `isw*()` function.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

SEE ALSO [iswalphabet\(3C\)](#), [setlocale\(3C\)](#), [wctype\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#), [standards\(5\)](#)

killpg(3C)

NAME	killpg – send signal to a process group						
SYNOPSIS	<pre>#include <signal.h> int killpg(pid_t pgrp, int sig);</pre>						
DESCRIPTION	<p>The <code>killpg()</code> function sends the signal <i>sig</i> to the process group <i>pgrp</i>. See <code>signal.h(3HEAD)</code> for a list of signals.</p> <p>The real or effective user ID of the sending process must match the real or saved set-user ID of the receiving process, unless the effective user ID of the sending process is the privileged user. A single exception is the signal <code>SIGCONT</code>, which may always be sent to any descendant of the current process.</p>						
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.						
ERRORS	The <code>killpg()</code> function will fail and no signal will be sent if:						
	<table><tr><td><code>EINVAL</code></td><td>The <i>sig</i> argument is not a valid signal number.</td></tr><tr><td><code>EPERM</code></td><td>The effective user ID of the sending process is not privileged user, and neither its real nor effective user ID matches the real or saved set-user ID of one or more of the target processes.</td></tr><tr><td><code>ESRCH</code></td><td>No processes were found in the specified process group.</td></tr></table>	<code>EINVAL</code>	The <i>sig</i> argument is not a valid signal number.	<code>EPERM</code>	The effective user ID of the sending process is not privileged user, and neither its real nor effective user ID matches the real or saved set-user ID of one or more of the target processes.	<code>ESRCH</code>	No processes were found in the specified process group.
<code>EINVAL</code>	The <i>sig</i> argument is not a valid signal number.						
<code>EPERM</code>	The effective user ID of the sending process is not privileged user, and neither its real nor effective user ID matches the real or saved set-user ID of one or more of the target processes.						
<code>ESRCH</code>	No processes were found in the specified process group.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>kill(2)</code> , <code>setpgrp(2)</code> , <code>sigaction(2)</code> , <code>signal.h(3HEAD)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME lckpwnf, ulckpwnf – manipulate shadow password database lock file

SYNOPSIS

```
#include <shadow.h>

int lckpwnf(void);

int ulckpwnf(void);
```

DESCRIPTION The lckpwnf() and ulckpwnf() functions enable modification access to the password databases through the lock file. A process first uses lckpwnf() to lock the lock file, thereby gaining exclusive rights to modify the /etc/passwd or /etc/shadow password database. See passwd(4) and shadow(4). Upon completing modifications, a process should release the lock on the lock file using ulckpwnf(). This mechanism prevents simultaneous modification of the password databases. The lock file, /etc/.pwd.lock, is used to coordinate modification access to the password databases /etc/passwd and /etc/shadow.

RETURN VALUES If lckpwnf() is successful in locking the file within 15 seconds, it returns 0. If unsuccessful (for example, /etc/.pwd.lock is already locked), it returns -1.

If ulckpwnf() is successful in unlocking the file /etc/.pwd.lock, it returns 0. If unsuccessful (for example, /etc/.pwd.lock is already unlocked), it returns -1.

USAGE These routines are for internal use only; compatibility is not guaranteed.

FILES

/etc/passwd	password database
/etc/shadow	shadow password database
/etc/.pwd.lock	lock file

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO getpwnam(3C), getspnam(3C), passwd(4), shadow(4), attributes(5)

lfmt(3C)

NAME	lfmt – display error message in standard format and pass to logging and monitoring services
SYNOPSIS	<pre>#include <pfmt.h> int lfmt(FILE *stream, long flags, char *format, ... /* arg* */);</pre>
DESCRIPTION	<p>The <code>lfmt()</code> function retrieves a format string from a locale-specific message database (unless <code>MM_NOGET</code> is specified) and uses it for <code>printf(3C)</code> style formatting of <i>args</i>. The output is displayed on <i>stream</i>. If <i>stream</i> is <code>NULL</code> no output is displayed.</p> <p>The <code>lfmt()</code> function encapsulates the output in the standard error message format (unless <code>MM_NOSTD</code> is specified, in which case the output is like that of <code>printf()</code>). It forwards its output to the logging and monitoring facility, even if <i>stream</i> is <code>NULL</code>. Optionally, <code>lfmt()</code> displays the output on the console with a date and time stamp.</p> <p>If the <code>printf()</code> format string is to be retrieved from a message database, the <i>format</i> argument must have the following structure:</p> <pre><catalog> : <msgnum> : <defmsg>.</pre> <p>If <code>MM_NOGET</code> is specified, only the <i><defmsg></i> field must be specified.</p> <p>The <i><catalog></i> field indicates the message database that contains the localized version of the format string. This field is limited to 14 characters selected from a set of all characters values, excluding the null character (<code>\0</code>) and the ASCII codes for slash (<code>/</code>) and colon (<code>:</code>).</p> <p>The <i><msgnum></i> field is a positive number that indicates the index of the string into the message database.</p> <p>If the catalog does not exist in the locale (specified by the last call to <code>setlocale(3C)</code> using the <code>LC_ALL</code> or <code>LC_MESSAGES</code> categories), or if the message number is out of bound, <code>lfmt()</code> will attempt to retrieve the message from the C locale. If this second retrieval fails, <code>lfmt()</code> uses the <i><defmsg></i> field of the <i>format</i> argument.</p> <p>If <i><catalog></i> is omitted, <code>lfmt()</code> will attempt to retrieve the string from the default catalog specified by the last call to <code>setcat(3C)</code>. In this case, the <i>format</i> argument has the following structure:</p> <pre>: <msgnum> : <defmsg>.</pre> <p>The <code>lfmt()</code> function will output the message</p> <pre>Message not found!!\n</pre> <p>as the format string if <i><catalog></i> is not a valid catalog name, if no catalog is specified (either explicitly or with <code>setcat()</code>), if <i><msgnum></i> is not a valid number, or if no message could be retrieved from the message databases and <i><defmsg></i> was omitted.</p>

The *flags* argument determines the type of output (whether the `format` should be interpreted as it is or be encapsulated in the standard message format) and the access to message catalogs to retrieve a localized version of `format`.

The *flags* argument is composed of several groups, and can take the following values (one from each group):

Output format control

MM_NOSTD	Do not use the standard message format but interpret <code>format</code> as a <code>printf()</code> format. Only <i>catalog access control flags</i> , <i>console display control</i> and <i>logging information</i> should be specified if MM_NOSTD is used; all other flags will be ignored.
MM_STD	Output using the standard message format (default value is 0).

Catalog access control

MM_NOGET	Do not retrieve a localized version of <code>format</code> . In this case, only the <code><defmsg></code> field of <code>format</code> is specified.
MM_GET	Retrieve a localized version of <code>format</code> from <code><catalog></code> , using <code><msgid></code> as the index and <code><defmsg></code> as the default message (default value is 0).

Severity (standard message format only)

MM_HALT	Generate a localized version of HALT, but donot halt the machine.
MM_ERROR	Generate a localized version of ERROR (default value is 0).
MM_WARNING	Generate a localized version of WARNING.
MM_INFO	Generate a localized version of INFO.

Additional severities can be defined with the [addsev\(3C\)](#) function, using number-string pairs with numeric values in the range [5-255]. The specified severity is formed by the bitwise OR operation of the numeric value and other *flags* arguments.

If the severity is not defined, `lfmt()` uses the string `SEV=N` where *N* is the integer severity value passed in *flags*.

Multiple severities passed in *flags* will not be detected as an error. Any combination of severities will be summed and the numeric value will cause the display of either a severity string (if defined) or the string `SEV=N` (if undefined).

Action

MM_ACTION	Specify an action message. Any severity value is superseded and replaced by a localized version of TO FIX.
-----------	--

lfmt(3C)

Console display control

- `MM_CONSOLE` Display the message to the console in addition to the specified *stream*.
- `MM_NOCONSOLE` Do not display the message to the console in addition to the specified *stream* (default value is 0).

Logging information

Major classification

Identify the source of the condition. Identifiers are: `MM_HARD` (hardware), `MM_SOFT` (software), and `MM_FIRM` (firmware).

Message source subclassification

Identify the type of software in which the problem is spotted. Identifiers are: `MM_APPL` (application), `MM_UTIL` (utility), and `MM_OPSYS` (operating system).

STANDARD ERROR MESSAGE FORMAT

The `lfmt()` function displays error messages in the following format:

label: *severity*: *text*

If no *label* was defined by a call to `setLabel(3C)`, the message is displayed in the format:

severity: *text*

If `lfmt()` is called twice to display an error message and a helpful *action* or recovery message, the output may appear as follows:

label: *severity*: *text*

label: TO FIX: *text*

RETURN VALUES

Upon successful completion, `lfmt()` returns the number of bytes transmitted. Otherwise, it returns a negative value:

- 1 Write the error to *stream*.
- 2 Cannot log and/or display at console.

USAGE

Since `lfmt()` uses `gettext(3C)`, it is recommended that `lfmt()` not be used.

EXAMPLES

EXAMPLE 1 The following example

```
setLabel("UX:test");
lfmt(stderr, MM_ERROR|MM_CONSOLE|MM_SOFT|MM_UTIL,
      "test:2:Cannot open file: %s\n", strerror(errno));
```

displays the message to `stderr` and to the console and makes it available for logging:

```
UX:test: ERROR: Cannot open file: No such file or directory
```

EXAMPLE 2 The following example

```
setLabel("UX:test");
lfmt(stderr, MM_INFO|MM_SOFT|MM_UTIL,
      "test:23:test facility is enabled\n");
```

EXAMPLE 2 The following example *(Continued)*

displays the message to `stderr` and makes it available for logging:

```
UX:test: INFO: test facility enabled
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `addsev(3C)`, `gettxt(3C)`, `pfmt(3C)`, `printf(3C)`, `setcat(3C)`, `setlabel(3C)`, `setlocale(3C)`, `attributes(5)`, `environ(5)`

localeconv(3C)

NAME	localeconv – get numeric formatting information
SYNOPSIS	<pre>#include <locale.h> struct lconv *localeconv(void);</pre>
DESCRIPTION	<p>The <code>localeconv()</code> function sets the components of an object with type <code>struct lconv</code> (defined in <code><locale.h></code>) with the values appropriate for the formatting of numeric quantities (monetary and otherwise) according to the rules of the current locale (see setlocale(3C)). The definition of <code>struct lconv</code> is given below (the values for the fields in the “C” locale are given in comments).</p> <pre>char *decimal_point; /* "." */ char *thousands_sep; /* "" (zero length string) */ char *grouping; /* "" */ char *int_curr_symbol; /* "" */ char *currency_symbol; /* "" */ char *mon_decimal_point; /* "" */ char *mon_thousands_sep; /* "" */ char *mon_grouping; /* "" */ char *positive_sign; /* "" */ char *negative_sign; /* "" */ char int_frac_digits; /* CHAR_MAX */ char frac_digits; /* CHAR_MAX */ char p_cs_precedes; /* CHAR_MAX */ char p_sep_by_space; /* CHAR_MAX */ char n_cs_precedes; /* CHAR_MAX */ char n_sep_by_space; /* CHAR_MAX */ char p_sign_posn; /* CHAR_MAX */ char n_sign_posn; /* CHAR_MAX */</pre> <p>The following members are also available to SUSv3–conforming applications. See standards(5)</p> <pre>char int_p_cs_precedes; /* CHAR_MAX */ char int_p_sep_by_space; /* CHAR_MAX */ char int_n_cs_precedes; /* CHAR_MAX */ char int_n_sep_by_space; /* CHAR_MAX */ char int_p_sign_posn; /* CHAR_MAX */ char int_n_sign_posn; /* CHAR_MAX */</pre> <p>The members of the structure with type <code>char *</code> are strings, any of which (except <code>decimal_point</code>) can point to a null string (“”), to indicate that the value is not available in the current locale or is of zero length. The members with type <code>char</code> are non-negative numbers, any of which can be <code>CHAR_MAX</code> (defined in the <code><limits.h></code> header) to indicate that the value is not available in the current locale. The members are the following:</p> <pre>char *decimal_point The decimal-point character used to format non-monetary quantities. char *thousands_sep The character used to separate groups of digits to the left of the decimal-point character in formatted non-monetary quantities.</pre>

`char *grouping`
A string whose elements taken as one-byte integer values indicate the size of each group of digits in formatted non-monetary quantities.

`char *int_curr_symbol`
The international currency symbol applicable to the current locale. The first three characters contain the alphabetic international currency symbol in accordance with those specified in the ISO 4217: 1995 standard. The fourth character (immediately preceding the null byte) is the character used to separate the international currency symbol from the monetary quantity.

`char *currency_symbol`
The local currency symbol applicable to the current locale.

`char *mon_decimal_point`
The decimal point used to format monetary quantities.

`char *mon_thousands_sep`
The separator for groups of digits to the left of the decimal point in formatted monetary quantities.

`char *mon_grouping`
A string whose elements taken as one-byte integer values indicate the size of each group of digits in formatted monetary quantities.

`char *positive_sign`
The string used to indicate a non-negative-valued formatted monetary quantity.

`char *negative_sign`
The string used to indicate a negative-valued formatted monetary quantity.

`char int_frac_digits`
The number of fractional digits (those to the right of the decimal point) to be displayed in an internationally formatted monetary quantity.

`char frac_digits`
The number of fractional digits (those to the right of the decimal point) to be displayed in a formatted monetary quantity.

`char p_cs_precedes`
Set to 1 or 0 if the `currency_symbol` respectively precedes or succeeds the value for a non-negative formatted monetary quantity.

`char p_sep_by_space`
Set to 0 if no space separates the `currency_symbol` or `int_curr_symbol` from the value for a non-negative formatted monetary quantity. Set to 1 if a space separates the symbol from the value; and set to 2 if a space separates the symbol and the sign string, if adjacent.

`char n_cs_precedes`
Set to 1 or 0 if the `currency_symbol` respectively precedes or succeeds the value for a negative formatted monetary quantity.

localeconv(3C)

`char n_sep_by_space`
Set to 0 if no space separates the `currency_symbol` or `int_curr_symbol` from the value for a negative formatted monetary quantity. Set to 1 if a space separates the symbol from the value; and set to 2 if a space separates the symbol and the sign string, if adjacent.

`char p_sign_posn`
Set to a value indicating the positioning of the `positive_sign` for a non-negative formatted monetary quantity.

`char n_sign_posn`
Set to a value indicating the positioning of the `negative_sign` for a negative formatted monetary quantity.

`char int_p_cs_precedes`
Set to 1 or 0 if the `int_curr_symbol` respectively precedes or succeeds the value for a non-negative internationally formatted monetary quantity.

`char int_n_cs_precedes`
Set to 1 or 0 if the `int_curr_symbol` respectively precedes or succeeds the value for a negative internationally formatted monetary quantity.

`char int_p_sep_by_space`
Set to a value indicating the separation of the `int_curr_symbol`, the sign string, and the value for a non-negative internationally formatted monetary quantity.

`char int_n_sep_by_space`
Set to a value indicating the separation of the `int_curr_symbol`, the sign string, and the value for a negative internationally formatted monetary quantity.

`char int_p_sign_posn`
Set to a value indicating the positioning of the `positive_sign` for a non-negative internationally formatted monetary quantity.

`char int_n_sign_posn`
Set to a value indicating the positioning of the `negative_sign` for a negative internationally formatted monetary quantity.

The elements of `grouping` and `mon_grouping` are interpreted according to the following:

<code>{CHAR_MAX}</code>	No further grouping is to be performed.
0	The previous element is to be repeatedly used for the remainder of the digits.
<i>other</i>	The integer value is the number of digits that comprise the current group. The next element is examined to determine the size of the next group of digits before the current group.

The values of `p_sep_by_space`, `n_sep_by_space`, `int_p_sep_by_space`, and `int_n_sep_by_space` are interpreted according to the following:

0	No space separates the currency symbol and value.
---	---

- 1 If the currency symbol and sign string are adjacent, a space separates them from the value; otherwise, a space separates the currency symbol from the value.
- 2 If the currency symbol and sign string are adjacent, a space separates them; otherwise, a space separates the sign string from the value.

In an SUSv3-conforming application, for `int_p_sep_by_space` and `int_n_sep_by_space`, the fourth character of `int_curr_symbol` is used instead of a space.

The values of `p_sign_posn`, `n_sign_posn`, `int_p_sign_posn`, and `int_n_sign_posn` are interpreted according to the following:

- 0 Parentheses surround the quantity and `currency_symbol` or `int_curr_symbol`.
- 1 The sign string precedes the quantity and `currency_symbol` or `int_curr_symbol`.
- 2 The sign string succeeds the quantity and `currency_symbol` or `int_curr_symbol`.
- 3 The sign string immediately precedes the `currency_symbol` or `int_curr_symbol`.
- 4 The sign string immediately succeeds the `currency_symbol` or `int_curr_symbol`.

RETURN VALUES

The `localeconv()` function returns a pointer to the filled-in object. The structure pointed to by the return value may be overwritten by a subsequent call to `localeconv()`.

EXAMPLES

EXAMPLE 1 Rules used by four countries to format monetary quantities.

The following table illustrates the rules used by four countries to format monetary quantities.

Country	Positive	Negative	International
Italy (IT)	L.1.234	-L.1.234	ITL.1.234
Netherlands (NE)	F 1.234,56	F -1.234,56	NLG 1.234,56
Norway (NO)	kr1.234,56	kr1.234,56-	NOK 1.234,56
Switzerland (SW)	SFrs.1,234.56	SFrs.1,234.56C	CHF 1,234.56

For these four countries, the respective values for the monetary members of the structure returned by `localeconv()` are as follows:

localeconv(3C)

	IT	NE	NO	SW
int_curr_symbol	"ITL."	"NLG "	"NOK "	"CHF "
currency_symbol	"L."	"F"	"kr"	"SFrs."
mon_decimal_point	""	","	","	":"
mon_thousands_sep	"."	."	."	","
mon_grouping	"\3"	"\3"	"\3"	"\3"
positive_sign	""	""	""	""
negative_sign	"_"	"_"	"_"	"C"
int_frac_digits	0	2	2	2
frac_digits	0	2	2	2
p_cs_precedes	1	1	1	1
p_sep_by_space	0	1	0	0
n_cs_precedes	1	1	1	1
n_sep_by_space	0	1	0	0
p_sign_posn	1	1	1	1
n_sign_posn	1	4	2	2
int_p_cs_precedes	1	1	1	1
int_n_cs_precedes	1	1	1	1
int_p_sep_by_space	0	0	0	0
int_n_sep_by_space	0	0	0	0
int_p_sign_posn	1	1	1	1
int_n_sign_posn	1	4	4	2

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

The `localeconv()` function can be used safely in multithreaded applications, as long as `setlocale(3C)` is not being called to change the locale.

localeconv(3C)

SEE ALSO [setlocale\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#), [standards\(5\)](#)

lockf(3C)

NAME	lockf – record locking on files
SYNOPSIS	<pre>#include <unistd.h> int lockf(int <i>fildes</i>, int <i>function</i>, off_t <i>size</i>);</pre>
DESCRIPTION	<p>The <code>lockf()</code> function allows sections of a file to be locked; advisory or mandatory write locks depending on the mode bits of the file (see <code>chmod(2)</code>). Calls to <code>lockf()</code> from other threads that attempt to lock the locked file section will either return an error value or be put to sleep until the resource becomes unlocked. All the locks for a process are removed when the process terminates. See <code>fcntl(2)</code> for more information about record locking.</p> <p>The <i>fildes</i> argument is an open file descriptor. The file descriptor must have <code>O_WRONLY</code> or <code>O_RDWR</code> permission in order to establish locks with this function call.</p> <p>The <i>function</i> argument is a control value that specifies the action to be taken. The permissible values for <i>function</i> are defined in <code><unistd.h></code> as follows:</p> <pre>#define F_ULOCK 0 /* unlock previously locked section */ #define F_LOCK 1 /* lock section for exclusive use */ #define F_TLOCK 2 /* test & lock section for exclusive use */ #define F_TEST 3 /* test section for other locks */</pre> <p>All other values of <i>function</i> are reserved for future extensions and will result in an error if not implemented.</p> <p><code>F_TEST</code> is used to detect if a lock by another process is present on the specified section. <code>F_LOCK</code> and <code>F_TLOCK</code> both lock a section of a file if the section is available. <code>F_ULOCK</code> removes locks from a section of the file.</p> <p>The <i>size</i> argument is the number of contiguous bytes to be locked or unlocked. The resource to be locked or unlocked starts at the current offset in the file and extends forward for a positive <i>size</i> and backward for a negative <i>size</i> (the preceding bytes up to but not including the current offset). If <i>size</i> is zero, the section from the current offset through the largest file offset is locked (that is, from the current offset through the present or any future end-of-file). An area need not be allocated to the file in order to be locked as such locks may exist past the end-of-file.</p> <p>The sections locked with <code>F_LOCK</code> or <code>F_TLOCK</code> may, in whole or in part, contain or be contained by a previously locked section for the same process. Locked sections will be unlocked starting at the point of the offset through <i>size</i> bytes or to the end of file if <i>size</i> is <code>(off_t) 0</code>. When this situation occurs, or if this situation occurs in adjacent sections, the sections are combined into a single section. If the request requires that a new element be added to the table of active locks and this table is already full, an error is returned, and the new section is not locked.</p> <p><code>F_LOCK</code> and <code>F_TLOCK</code> requests differ only by the action taken if the resource is not available. <code>F_LOCK</code> blocks the calling thread until the resource is available. <code>F_TLOCK</code> causes the function to return <code>-1</code> and set <code>errno</code> to <code>EAGAIN</code> if the section is already locked by another process.</p>

File locks are released on first close by the locking process of any file descriptor for the file.

`F_ULOCK` requests may, in whole or in part, release one or more locked sections controlled by the process. When sections are not fully released, the remaining sections are still locked by the process. Releasing the center section of a locked section requires an additional element in the table of active locks. If this table is full, an `errno` is set to `EDEADLK` and the requested section is not released.

An `F_ULOCK` request in which `size` is non-zero and the offset of the last byte of the requested section is the maximum value for an object of type `off_t`, when the process has an existing lock in which `size` is 0 and which includes the last byte of the requested section, will be treated as a request to unlock from the start of the requested section with a `size` equal to 0. Otherwise, an `F_ULOCK` request will attempt to unlock only the requested section.

A potential for deadlock occurs if the threads of a process controlling a locked resource is put to sleep by requesting another process's locked resource. Thus calls to `lockf()` or `fcntl(2)` scan for a deadlock prior to sleeping on a locked resource. An error return is made if sleeping on the locked resource would cause a deadlock.

Sleeping on a resource is interrupted with any signal. The `alarm(2)` function may be used to provide a timeout facility in applications that require this facility.

RETURN VALUES

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

ERRORS

The `lockf()` function will fail if:

<code>EBADF</code>	The <i>fildev</i> argument is not a valid open file descriptor; or function is <code>F_LOCK</code> or <code>F_TLOCK</code> and <i>fildev</i> is not a valid file descriptor open for writing.
<code>EACCES</code> or <code>EAGAIN</code>	The function argument is <code>F_TLOCK</code> or <code>F_TEST</code> and the section is already locked by another process.
<code>EDEADLK</code>	The function argument is <code>F_LOCK</code> and a deadlock is detected.
<code>EINTR</code>	A signal was caught during execution of the function.
<code>ECOMM</code>	The <i>fildev</i> argument is on a remote machine and the link to that machine is no longer active.
<code>EINVAL</code>	The function argument is not one of <code>F_LOCK</code> , <code>F_TLOCK</code> , <code>F_TEST</code> , or <code>F_ULOCK</code> ; or <code>size</code> plus the current file offset is less than 0.
<code>EOVERFLOW</code>	The offset of the first, or if <code>size</code> is not 0 then the last, byte in the requested section cannot be represented correctly in an object of type <code>off_t</code> .

The `lockf()` function may fail if:

lockf(3C)

EAGAIN	The function argument is <code>F_LOCK</code> or <code>F_TLOCK</code> and the file is mapped with <code>mmap(2)</code> .
EDEADLK or ENOLCK	The function argument is <code>F_LOCK</code> , <code>F_TLOCK</code> , or <code>F_ULOCK</code> and the request would cause the number of locks to exceed a system-imposed limit.
EOPNOTSUPP or EINVAL	The locking of files of the type indicated by the <i>fildev</i> argument is not supported.

USAGE Record-locking should not be used in combination with the `fopen(3C)`, `fread(3C)`, `fwrite(3C)` and other `stdio` functions. Instead, the more primitive, non-buffered functions (such as `open(2)`) should be used. Unexpected results may occur in processes that do buffering in the user address space. The process may later read/write data which is/was locked. The `stdio` functions are the most common source of unexpected buffering.

The `alarm(2)` function may be used to provide a timeout facility in applications requiring it.

The `lockf()` 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
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `intro(2)`, `alarm(2)`, `chmod(2)`, `close(2)`, `creat(2)`, `fcntl(2)`, `mmap(2)`, `open(2)`, `read(2)`, `write(2)`, `attributes(5)`, `lf64(5)`, `standards(5)`

NAME | `_longjmp, _setjmp` – non-local goto

SYNOPSIS | `#include <setjmp.h>`
`void _longjmp(jmp_buf env, int val);`
`int _setjmp(jmp_buf env);`

DESCRIPTION | The `_longjmp()` and `_setjmp()` functions are identical to `longjmp(3C)` and `setjmp(3C)`, respectively, with the additional restriction that `_longjmp()` and `_setjmp()` do not manipulate the signal mask.

If `_longjmp()` is called even though `env` was never initialized by a call to `_setjmp()`, or when the last such call was in a function that has since returned, the results are undefined.

RETURN VALUES | Refer to `longjmp(3C)` and `setjmp(3C)`.

ERRORS | No errors are defined.

USAGE | If `_longjmp()` is executed and the environment in which `_setjmp()` was executed no longer exists, errors can occur. The conditions under which the environment of the `_setjmp()` no longer exists include exiting the function that contains the `_setjmp()` call, and exiting an inner block with temporary storage. This condition might not be detectable, in which case the `_longjmp()` occurs and, if the environment no longer exists, the contents of the temporary storage of an inner block are unpredictable. This condition might also cause unexpected process termination. If the function has returned, the results are undefined.

Passing `longjmp()` a pointer to a buffer not created by `setjmp()`, passing `_longjmp()` a pointer to a buffer not created by `_setjmp()`, passing `siglongjmp(3C)` a pointer to a buffer not created by `sigsetjmp(3C)` or passing any of these three functions a buffer that has been modified by the user can cause all the problems listed above, and more.

The `_longjmp()` and `_setjmp()` functions are included to support programs written to historical system interfaces. New applications should use `siglongjmp(3C)` and `sigsetjmp(3C)` respectively.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

SEE ALSO | `longjmp(3C)`, `setjmp(3C)`, `siglongjmp(3C)`, `sigsetjmp(3C)`, `attributes(5)`, `standards(5)`

lsearch(3C)

NAME	lsearch, lfind – linear search and update
SYNOPSIS	<pre>#include <search.h> void *lsearch(const void *key, void *base, size_t *nelp, size_t width, int (*compar) (const void *, const void *)); void *lfind(const void *key, const void *base, size_t *nelp, size_t width, int (*compar) (const void *, const void *));</pre>
DESCRIPTION	<p>The <code>lsearch()</code> function is a linear search routine generalized from Knuth (6.1) Algorithm S. (see <i>The Art of Computer Programming, Volume 3, Section 6.1, by Donald E. Knuth.</i>) It returns a pointer to a table indicating where a datum can be found. If the datum does not occur, it is added at the end of the table. The <i>key</i> argument points to the datum to be sought in the table. The <i>base</i> argument points to the first element in the table. The <i>nelp</i> argument points to an integer containing the current number of elements in the table. The integer is incremented if the datum is added to the table. The <i>width</i> argument is the size of an element in bytes. The <i>compar</i> argument is a pointer to the comparison function that the user must supply (<code>strcmp(3C)</code> for example). It is called with two arguments that point to the elements being compared. The function must return zero if the elements are equal and non-zero otherwise.</p> <p>The <code>lfind()</code> function is the same as <code>lsearch()</code> except that if the datum is not found, it is not added to the table. Instead, a null pointer is returned.</p> <p>It is important to note the following:</p> <ul style="list-style-type: none">■ The pointers to the key and the element at the base of the table can be pointers to any type.■ The comparison function need not compare every byte, so arbitrary data can be contained in the elements in addition to the values being compared.■ The value returned should be cast into type pointer-to-element.
RETURN VALUES	If the searched-for datum is found, both <code>lsearch()</code> and <code>lfind()</code> return a pointer to it. Otherwise, <code>lfind()</code> returns <code>NULL</code> and <code>lsearch()</code> returns a pointer to the newly added element.
USAGE	Undefined results can occur if there is not enough room in the table to add a new item.
EXAMPLES	<p>EXAMPLE 1 A sample code using the <code>lsearch()</code> function.</p> <p>This program will read in less than <code>TABSIZE</code> strings of length less than <code>ELSIZE</code> and store them in a table, eliminating duplicates, and then will print each entry.</p> <pre>#include <search.h> #include <string.h> #include <stdlib.h> #include <stdio.h></pre>

EXAMPLE 1 A sample code using the `lsearch()` function. (Continued)

```
#define TABSIZE 50
#define ELSIZE 120

main()
{
    char line[ELSIZE];           /* buffer to hold input string */
    char tab[TABSIZE][ELSIZE];  /* table of strings */
    size_t nel = 0;             /* number of entries in tab */
    int i;

    while (fgets(line, ELSIZE, stdin) != NULL &&
           nel < TABSIZE)
        (void) lsearch(line, tab, &nel, ELSIZE, mycmp);
    for( i = 0; i < nel; i++ )
        (void) fputs(tab[i], stdout);
    return 0;
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `bsearch(3C)`, `hsearch(3C)`, `string(3C)`, `tsearch(3C)`, `attributes(5)`, `standards(5)`

The Art of Computer Programming, Volume 3, Sorting and Searching by Donald E. Knuth, published by Addison-Wesley Publishing Company, 1973.

madvise(3C)

NAME	madvise – provide advice to VM system
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/mman.h> int madvise(caddr_t <i>addr</i>, size_t <i>len</i>, int <i>advice</i>);</pre>
DESCRIPTION	<p>The <code>madvise()</code> function advises the kernel that a region of user mapped memory in the range [<i>addr</i>, <i>addr</i> + <i>len</i>) will be accessed following a type of pattern. The kernel uses this information to optimize the procedure for manipulating and maintaining the resources associated with the specified mapping range.</p> <p>Values for <i>advice</i> are defined in <code><sys/mman.h></code> as:</p> <pre>#define MADV_NORMAL 0x0 /* No further special treatment */ #define MADV_RANDOM 0x1 /* Expect random page references */ #define MADV_SEQUENTIAL 0x2 /* Expect sequential page references */ #define MADV_WILLNEED 0x3 /* Will need these pages */ #define MADV_DONTNEED 0x4 /* Don't need these pages */ #define MADV_FREE 0x5 /* Contents can be freed */ #define MADV_ACCESS_DEFAULT 0x6 /* default access */ #define MADV_ACCESS_LWP 0x7 /* next LWP to access heavily */ #define MADV_ACCESS_MANY 0x8 /* many processes to access heavily */</pre> <p>MADV_NORMAL The default system characteristic where accessing memory within the address range causes the system to read data from the mapped file. The kernel reads all data from files into pages which are retained for a period of time as a “cache.” System pages can be a scarce resource, so the kernel steals pages from other mappings when needed. This is a likely occurrence, but adversely affects system performance only if a large amount of memory is accessed.</p> <p>MADV_RANDOM Tells the kernel to read in a minimum amount of data from a mapped file on any single particular access. If MADV_NORMAL is in effect when an address of a mapped file is accessed, the system tries to read in as much data from the file as reasonable, in anticipation of other accesses within a certain locality.</p> <p>MADV_SEQUENTIAL Tells the system that addresses in this range are likely to be accessed only once, so the system will free the resources mapping the address range as quickly as possible.</p> <p>MADV_WILLNEED Tells the system that a certain address range is definitely needed so the kernel will start reading the specified range into memory. This can benefit programs wanting to minimize the time needed to access memory the first time, as the kernel would need to read in from the file.</p>

MADV_DONTNEED	Tells the kernel that the specified address range is no longer needed, so the system starts to free the resources associated with the address range.
MADV_FREE	Tells the kernel that contents in the specified address range are no longer important and the range will be overwritten. When there is demand for memory, the system will free pages associated with the specified address range. In this instance, the next time a page in the address range is referenced, it will contain all zeroes. Otherwise, it will contain the data that was there prior to the MADV_FREE call. References made to the address range will not make the system read from backing store (swap space) until the page is modified again. This value cannot be used on mappings that have underlying file objects.
MADV_ACCESS_LWP	Tells the kernel that the next LWP to touch the specified address range will access it most heavily, so the kernel should try to allocate the memory and other resources for this range and the LWP accordingly.
MADV_ACCESS_MANY	Tells the kernel that many processes and/or LWPs will access the specified address range randomly across the machine, so the kernel should try to allocate the memory and other resources for this range accordingly.
MADV_ACCESS_DEFAULT	Resets the kernel's expectation for how the specified range will be accessed to the default.

The `madvise()` function should be used by applications with specific knowledge of their access patterns over a memory object, such as a mapped file, to increase system performance.

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

ERRORS	EAGAIN	Some or all mappings in the address range [<i>addr</i> , <i>addr + len</i>) are locked for I/O.
	EBUSY	Some or all of the addresses in the range [<i>addr</i> , <i>addr + len</i>) are locked and <code>MS_SYNC</code> with the <code>MS_INVALIDATE</code> option is specified.
	EFAULT	Some or all of the addresses in the specified range could not be read into memory from the underlying object when performing <code>MADV_WILLNEED</code> .

madvise(3C)

- EINVAL The *addr* argument is not a multiple of the page size as returned by [sysconf\(3C\)](#), the length of the specified address range is equal to 0, or the *advice* argument was invalid.
- EIO An I/O error occurred while reading from or writing to the file system.
- ENOMEM Addresses in the range [*addr*, *addr + len*) are outside the valid range for the address space of a process, or specify one or more pages that are not mapped.
- ESTALE Stale NFS file handle.

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

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

SEE ALSO [meminfo\(2\)](#), [mmap\(2\)](#), [sysconf\(3C\)](#), [attributes\(5\)](#)

NAME	makecontext, swapcontext – manipulate user contexts
SYNOPSIS	<pre>#include <ucontext.h> void makecontext(ucontext_t *ucp, void (*func)(), int argc, ...); int swapcontext(ucontext_t *restrict oucp, const ucontext_t *restrict ucp);</pre>
DESCRIPTION	<p>The <code>makecontext()</code> function modifies the context specified by <code>ucp</code>, which has been initialized using <code>getcontext(2)</code>. When this context is resumed using <code>swapcontext()</code> or <code>setcontext(2)</code>, execution continues by calling the function <code>func</code>, passing it the arguments that follow <code>argc</code> in the <code>makecontext()</code> call. The value of <code>argc</code> must match the number of pointer-sized integer arguments passed to <code>func</code>, otherwise the behavior is undefined.</p> <p>Before a call is made to <code>makecontext()</code>, the context being modified should have a stack allocated for it. The stack is assigned to the context by initializing the <code>uc_stack</code> member.</p> <p>The <code>uc_link</code> member is used to determine the context that will be resumed when the context being modified by <code>makecontext()</code> returns. The <code>uc_link</code> member should be initialized prior to the call to <code>makecontext()</code>. If the <code>uc_link</code> member is initialized to <code>NULL</code>, the thread executing <code>func</code> will exit when <code>func</code> returns. See pthread_exit(3C).</p> <p>The <code>swapcontext()</code> function saves the current context in the context structure pointed to by <code>oucp</code> and sets the context to the context structure pointed to by <code>ucp</code>.</p> <p>If the <code>ucp</code> or <code>oucp</code> argument points to an invalid address, the behavior is undefined and <code>errno</code> may be set to <code>EFAULT</code>.</p>
RETURN VALUES	On successful completion, <code>swapcontext()</code> returns 0. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The <code>swapcontext()</code> function will fail if:</p> <p>ENOMEM The <code>ucp</code> argument does not have enough stack left to complete the operation.</p> <p>The <code>swapcontext()</code> function may fail if:</p> <p>EFAULT The <code>ucp</code> or <code>oucp</code> argument points to an invalid address.</p>
EXAMPLES	<p>EXAMPLE 1 Alternate execution context on a stack whose memory was allocated using <code>mmap(2)</code>.</p> <pre>#include <stdio.h> #include <ucontext.h> #include <sys/mman.h> void assign(long a, int *b) {</pre>

makecontext(3C)

EXAMPLE 1 Alternate execution context on a stack whose memory was allocated using `mmap(2)`. (Continued)

```
        *b = (int)a;
    }

    int
    main(int argc, char **argv)
    {
        ucontext_t uc, back;
        size_t sz = 0x10000;
        int value = 0;

        getcontext(&uc);

        uc.uc_stack.ss_sp = mmap(0, sz,
            PROT_READ | PROT_WRITE | PROT_EXEC,
            MAP_PRIVATE | MAP_ANON, -1, 0);
        uc.uc_stack.ss_size = sz;
        uc.uc_stack.ss_flags = 0;

        uc.uc_link = &back;

        makecontext(&uc, assign, 2, 100L, &value);
        swapcontext(&back, &uc);

        printf("done %d\n", value);

        return (0);
    }
```

USAGE These functions are useful for implementing user-level context switching between multiple threads of control within a process (co-processing). More effective multiple threads of control can be obtained by using native support for multithreading. See `threads(5)`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `mmap(2)`, `getcontext(2)`, `sigaction(2)`, `sigprocmask(2)`, `pthread_exit(3C)`, `ucontext.h(3HEAD)`, `attributes(5)`, `standards(5)`, `threads(5)`

NOTES The semantics of the `uc_stack` member of the `ucontext_t` structure have changed as they apply to inputs to `makecontext()`. Prior to Solaris 10, the `ss_sp` member of the `uc_stack` structure represented the high memory address of the area reserved for the stack. The `ss_sp` member now represents the base (low memory address), in keeping with other uses of `ss_sp`.

This change in the meaning of `ss_sp` is now the default behavior. The `-D__MAKECONTEXT_V2_SOURCE` compilation flag used in Solaris 9 update releases to access this behavior is obsolete.

Binary compatibility has been preserved with releases prior to Solaris 10. Before recompiling, applications that use `makecontext()` must be updated to reflect this behavior change. The example below demonstrates a typical change that must be applied:

```
--- example1_s9.c      Thu Oct  3 11:58:17 2002
+++ example1.c        Thu Jun 27 13:28:16 2002
@@ -27,12 +27,9 @@
     uc.uc_stack.ss_sp = mmap(0, sz,
        PROT_READ | PROT_WRITE | PROT_EXEC,
        MAP_PRIVATE | MAP_ANON, -1, 0);
-   uc.uc_stack.ss_sp = (char *)uc.uc_stack.ss_sp + sz - 8;
   uc.uc_stack.ss_size = sz;
   uc.uc_stack.ss_flags = 0;

   uc.uc_link = &back

   makecontext(&uc, assign, 2, 100L, &value);
```

makedev(3C)

NAME makedev, major, minor – manage a device number

SYNOPSIS

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

dev_t makedev(major_t maj, minor_t min) ;
major_t major(dev_t device) ;
minor_t minor(dev_t device) ;
```

DESCRIPTION

The `makedev()` function returns a formatted device number on success and `NODEV` on failure. The *maj* argument is the major number. The *min* argument is the minor number. The `makedev()` function can be used to create a device number for input to `mknod(2)`.

The `major()` function returns the major number component from *device*.

The `minor()` function returns the minor number component from *device*.

RETURN VALUES

Upon successful completion, `makedev()` returns a formatted device number. Otherwise, `NODEV` is returned and `errno` is set to indicate the error.

ERRORS

The `makedev()` function will fail if:

`EINVAL` One or both of the arguments *maj* and *min* is too large, or the *device* number created from *maj* and *min* is `NODEV`.

The `major()` function will fail if:

`EINVAL` The *device* argument is `NODEV`, or the major number component of *device* is too large.

The `minor()` function will fail if:

`EINVAL` The *device* argument is `NODEV`.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `mknod(2)`, `stat(2)`, `attributes(5)`

NAME	malloc, calloc, free, memalign, realloc, valloc, alloca – memory allocator
SYNOPSIS	<pre>#include <stdlib.h> void *malloc(size_t size); void *calloc(size_t nelem, size_t elsize); void free(void *ptr); void *memalign(size_t alignment, size_t size); void *realloc(void *ptr, size_t size); void *valloc(size_t size); #include <alloca.h> void *alloca(size_t size);</pre>
DESCRIPTION	<p>The <code>malloc()</code> and <code>free()</code> functions provide a simple, general-purpose memory allocation package. The <code>malloc()</code> function returns a pointer to a block of at least <i>size</i> bytes suitably aligned for any use. If the space assigned by <code>malloc()</code> is overrun, the results are undefined.</p> <p>The argument to <code>free()</code> is a pointer to a block previously allocated by <code>malloc()</code>, <code>calloc()</code>, or <code>realloc()</code>. After <code>free()</code> is executed, this space is made available for further allocation by the application, though not returned to the system. Memory is returned to the system only upon termination of the application. If <i>ptr</i> is a null pointer, no action occurs. If a random number is passed to <code>free()</code>, the results are undefined.</p> <p>The <code>calloc()</code> function allocates space for an array of <i>nelem</i> elements of size <i>elsize</i>. The space is initialized to zeros.</p> <p>The <code>memalign()</code> function allocates <i>size</i> bytes on a specified alignment boundary and returns a pointer to the allocated block. The value of the returned address is guaranteed to be an even multiple of <i>alignment</i>. The value of <i>alignment</i> must be a power of two and must be greater than or equal to the size of a word.</p> <p>The <code>realloc()</code> function changes the size of the block pointed to by <i>ptr</i> to <i>size</i> bytes and returns a pointer to the (possibly moved) block. The contents will be unchanged up to the lesser of the new and old sizes. If the new size of the block requires movement of the block, the space for the previous instantiation of the block is freed. If the new size is larger, the contents of the newly allocated portion of the block are unspecified. If <i>ptr</i> is NULL, <code>realloc()</code> behaves like <code>malloc()</code> for the specified size. If <i>size</i> is 0 and <i>ptr</i> is not a null pointer, the space pointed to is freed.</p> <p>The <code>valloc()</code> function has the same effect as <code>malloc()</code>, except that the allocated memory will be aligned to a multiple of the value returned by <code>sysconf(_SC_PAGESIZE)</code>.</p>

malloc(3C)

The `alloca()` function allocates *size* bytes of space in the stack frame of the caller, and returns a pointer to the allocated block. This temporary space is automatically freed when the caller returns. If the allocated block is beyond the current stack limit, the resulting behavior is undefined.

RETURN VALUES Upon successful completion, each of the allocation functions returns a pointer to space suitably aligned (after possible pointer coercion) for storage of any type of object.

If there is no available memory, `malloc()`, `realloc()`, `memalign()`, `valloc()`, and `calloc()` return a null pointer. When `realloc()` is called with *size* > 0 and returns `NULL`, the block pointed to by *ptr* is left intact. If *size*, *nelem*, or *elsize* is 0, either a null pointer or a unique pointer that can be passed to `free()` is returned.

If `malloc()`, `calloc()`, or `realloc()` returns unsuccessfully, `errno` will be set to indicate the error. The `free()` function does not set `errno`.

ERRORS The `malloc()`, `calloc()`, and `realloc()` functions will fail if:

- | | |
|---------------------|--|
| <code>ENOMEM</code> | The physical limits of the system are exceeded by <i>size</i> bytes of memory which cannot be allocated. |
| <code>EAGAIN</code> | There is not enough memory available to allocate <i>size</i> bytes of memory; but the application could try again later. |

USAGE Portable applications should avoid using `valloc()` but should instead use `malloc()` or `mmap(2)`. On systems with a large page size, the number of successful `valloc()` operations might be 0.

Comparative features of `malloc(3C)`, `bsdmalloc(3MALLOC)`, and `malloc(3MALLOC)` are as follows:

- The `bsdmalloc(3MALLOC)` routines afford better performance, but are space-inefficient.
- The `malloc(3MALLOC)` routines are space-efficient, but have slower performance.
- The standard, fully SCD-compliant `malloc` routines are a trade-off between performance and space-efficiency.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	<code>malloc()</code> , <code>calloc()</code> , <code>free()</code> , <code>realloc()</code> , <code>valloc()</code> are Standard; <code>memalign()</code> and <code>alloca()</code> are Stable.
MT-Level	Safe

SEE ALSO `brk(2)`, `getrlimit(2)`, `bsdmalloc(3MALLOC)`, `malloc(3MALLOC)`, `mapmalloc(3MALLOC)`, `watchmalloc(3MALLOC)`, `attributes(5)`

WARNINGS Undefined results will occur if the size requested for a block of memory exceeds the maximum size of a process's heap, which can be obtained with `getrlimit(2)`

The `alloca()` function is machine-, compiler-, and most of all, system-dependent. Its use is strongly discouraged.

malloc(3MALLOC)

NAME	malloc, free, realloc, calloc, malloc, mallinfo – memory allocator				
SYNOPSIS	<pre>cc [flag ...] file ... -lmalloc [library ...] #include <stdlib.h> void *malloc(size_t size); void free(void *ptr); void *realloc(void *ptr, size_t size); void *calloc(size_t nelem, size_t elsize); #include <malloc.h> int mallopt(int cmd, int value); struct mallinfo mallinfo(void);</pre>				
DESCRIPTION	<p>The malloc() and free() functions provide a simple general-purpose memory allocation package.</p> <p>The malloc() function returns a pointer to a block of at least <i>size</i> bytes suitably aligned for any use.</p> <p>The argument to free() is a pointer to a block previously allocated by malloc(). After free() is performed, this space is made available for further allocation, and its contents have been destroyed. See mallopt() below for a way to change this behavior. If <i>ptr</i> is a null pointer, no action occurs.</p> <p>Undefined results occur if the space assigned by malloc() is overrun or if some random number is handed to free().</p> <p>The realloc() function changes the size of the block pointed to by <i>ptr</i> to <i>size</i> bytes and returns a pointer to the (possibly moved) block. The contents will be unchanged up to the lesser of the new and old sizes. If the new size of the block requires movement of the block, the space for the previous instantiation of the block is freed. If the new size is larger, the contents of the newly allocated portion of the block are unspecified. If <i>ptr</i> is NULL, realloc() behaves like malloc() for the specified size. If <i>size</i> is 0 and <i>ptr</i> is not a null pointer, the space pointed to is freed.</p> <p>The calloc() function allocates space for an array of <i>nelem</i> elements of size <i>elsize</i>. The space is initialized to zeros.</p> <p>The mallopt() function provides for control over the allocation algorithm. The available values for <i>cmd</i> are:</p> <table><tr><td>M_MXFAST</td><td>Set <i>maxfast</i> to <i>value</i>. The algorithm allocates all blocks below the size of <i>maxfast</i> in large groups and then does them out very quickly. The default value for <i>maxfast</i> is 24.</td></tr><tr><td>M_NLBLKS</td><td>Set <i>numlblks</i> to <i>value</i>. The above mentioned “large groups” each contain <i>numlblks</i> blocks. <i>numlblks</i> must be greater than 0. The default value for <i>numlblks</i> is 100.</td></tr></table>	M_MXFAST	Set <i>maxfast</i> to <i>value</i> . The algorithm allocates all blocks below the size of <i>maxfast</i> in large groups and then does them out very quickly. The default value for <i>maxfast</i> is 24.	M_NLBLKS	Set <i>numlblks</i> to <i>value</i> . The above mentioned “large groups” each contain <i>numlblks</i> blocks. <i>numlblks</i> must be greater than 0. The default value for <i>numlblks</i> is 100.
M_MXFAST	Set <i>maxfast</i> to <i>value</i> . The algorithm allocates all blocks below the size of <i>maxfast</i> in large groups and then does them out very quickly. The default value for <i>maxfast</i> is 24.				
M_NLBLKS	Set <i>numlblks</i> to <i>value</i> . The above mentioned “large groups” each contain <i>numlblks</i> blocks. <i>numlblks</i> must be greater than 0. The default value for <i>numlblks</i> is 100.				

M_GRAIN	Set <i>grain</i> to <i>value</i> . The sizes of all blocks smaller than <i>maxfast</i> are considered to be rounded up to the nearest multiple of <i>grain</i> . <i>grain</i> must be greater than 0. The default value of <i>grain</i> is the smallest number of bytes that will allow alignment of any data type. Value will be rounded up to a multiple of the default when <i>grain</i> is set.
M_KEEP	Preserve data in a freed block until the next <code>malloc()</code> , <code>realloc()</code> , or <code>calloc()</code> . This option is provided only for compatibility with the old version of <code>malloc()</code> , and it is not recommended.

These values are defined in the `<malloc.h>` header.

The `mallocopt()` function can be called repeatedly, but cannot be called after the first small block is allocated.

The `mallinfo()` function provides instrumentation describing space usage. It returns the `mallinfo` structure with the following members:

```

unsigned long arena;      /* total space in arena */
unsigned long ordblks;   /* number of ordinary blocks */
unsigned long smlbks;    /* number of small blocks */
unsigned long hblkhd;    /* space in holding block headers */
unsigned long hblks;     /* number of holding blocks */
unsigned long usmlbks;   /* space in small blocks in use */
unsigned long fsmblks;   /* space in free small blocks */
unsigned long uordblks;  /* space in ordinary blocks in use */
unsigned long fordblks;  /* space in free ordinary blocks */
unsigned long keepcost;  /* space penalty if keep option */
                        /* is used */

```

The `mallinfo` structure is defined in the `<malloc.h>` header.

Each of the allocation routines returns a pointer to space suitably aligned (after possible pointer coercion) for storage of any type of object.

RETURN VALUES

The `malloc()`, `realloc()`, and `calloc()` functions return a null pointer if there is not enough available memory. When `realloc()` returns `NULL`, the block pointed to by *ptr* is left intact. If *size*, *nelem*, or *elsize* is 0, either a null pointer or a unique pointer that can be passed to `free()` is returned. If `mallocopt()` is called after any allocation or if *cmd* or *value* are invalid, a non-zero value is returned. Otherwise, it returns 0.

ERRORS

If `malloc()`, `calloc()`, or `realloc()` returns unsuccessfully, `errno` is set to indicate the error:

ENOMEM	<i>size</i> bytes of memory exceeds the physical limits of your system, and cannot be allocated.
EAGAIN	There is not enough memory available at this point in time to allocate <i>size</i> bytes of memory; but the application could try again later.

malloc(3MALLOC)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `brk(2)`, `bsdmalloc(3MALLOC)`, `libbmtmalloc(3LIB)`, `malloc(3C)`, `mapmalloc(3MALLOC)`, `mtmalloc(3MALLOC)`, `watchmalloc(3MALLOC)`, `attributes(5)`

NOTES Note that unlike `malloc(3C)`, this package does not preserve the contents of a block when it is freed, unless the `M_KEEP` option of `mallopt()` is used.

Undocumented features of `malloc(3C)` have not been duplicated.

Function prototypes for `malloc()`, `realloc()`, `calloc()`, and `free()` are also defined in the `<malloc.h>` header for compatibility with old applications. New applications should include `<stdlib.h>` to access the prototypes for these functions. Comparative Features of these `malloc` routines, `bsdmalloc(3MALLOC)`, and `malloc(3C)`

- These `malloc` routines are space-efficient but have slower performance.
- The `bsdmalloc(3MALLOC)` routines afford better performance but are space-inefficient.
- The standard, fully SCD-compliant `malloc(3C)` routines are a trade-off between performance and space-efficiency.

The `free()` function does not set `errno`.

NAME	mapmalloc – memory allocator
SYNOPSIS	<pre>cc [<i>flag ...</i>] <i>file ...</i> -lmapmalloc [<i>library ...</i>] #include <stdlib.h> void *malloc(size_t <i>size</i>); void *calloc(size_t <i>nelem</i>, size_t <i>elsize</i>); void free(void * <i>ptr</i>); void *realloc(void *<i>ptr</i>, size_t <i>size</i>);</pre>
DESCRIPTION	<p>The collection of malloc functions in this library use mmap(2) instead of sbrk(2) for acquiring new heap space. The functions in this library are intended to be used only if necessary, when applications must call sbrk(), but need to call other library routines that might call malloc. The algorithms used by these functions are not sophisticated. There is no reclaiming of memory.</p> <p>The malloc() and free() functions provide a simple general-purpose memory allocation package.</p> <p>The malloc() function returns a pointer to a block of at least <i>size</i> bytes suitably aligned for any use.</p> <p>The argument to free() is a pointer to a block previously allocated by malloc(), calloc() or realloc(). If <i>ptr</i> is a NULL pointer, no action occurs.</p> <p>Undefined results will occur if the space assigned by malloc() is overrun or if some random number is handed to free().</p> <p>The calloc() function allocates space for an array of <i>nelem</i> elements of size <i>elsize</i>. The space is initialized to zeros.</p> <p>The realloc() function changes the size of the block pointed to by <i>ptr</i> to <i>size</i> bytes and returns a pointer to the (possibly moved) block. The contents will be unchanged up to the lesser of the new and old sizes. If the new size of the block requires movement of the block, the space for the previous instantiation of the block is freed. If the new size is larger, the contents of the newly allocated portion of the block are unspecified. If <i>ptr</i> is NULL, realloc() behaves like malloc() for the specified size. If <i>size</i> is 0 and <i>ptr</i> is not a null pointer, the space pointed to is freed.</p> <p>Each of the allocation functions returns a pointer to space suitably aligned (after possible pointer coercion) for storage of any type of object.</p> <p>The malloc() and realloc() functions will fail if there is not enough available memory.</p> <p>Entry points for malloc_debug(), mallocmap(), mallopt(), mallinfo(), memalign(), and valloc() are empty routines, and are provided only to protect the user from mixing malloc() functions from different implementations.</p>

mapmalloc(3MALLOC)

RETURN VALUES If there is no available memory, `malloc()`, `realloc()`, and `calloc()` return a null pointer. When `realloc()` returns `NULL`, the block pointed to by `ptr` is left intact. If `size`, `nelem`, or `elsize` is 0, a unique pointer to the arena is returned.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `brk(2)`, `getrlimit(2)`, `mmap(2)`, `realloc(3C)`, `malloc(3MALLOC)`, `attributes(5)`

NAME	mblen – get number of bytes in a character								
SYNOPSIS	<pre>#include <stdlib.h> int mblen(const char *s, size_t n);</pre>								
DESCRIPTION	<p>If <i>s</i> is not a null pointer, <code>mblen()</code> determines the number of bytes constituting the character pointed to by <i>s</i>. It is equivalent to:</p> <pre>mbtowc((wchar_t *)0, s, n);</pre> <p>A call with <i>s</i> as a null pointer causes this function to return 0. The behavior of this function is affected by the <code>LC_CTYPE</code> category of the current locale.</p>								
RETURN VALUES	<p>If <i>s</i> is a null pointer, <code>mblen()</code> returns 0. If <i>s</i> is not a null pointer, <code>mblen()</code> returns 0 (if <i>s</i> points to the null byte), the number of bytes that constitute the character (if the next <i>n</i> or fewer bytes form a valid character), or -1 (if they do not form a valid character) and may set <code>errno</code> to indicate the error. In no case will the value returned be greater than <i>n</i> or the value of the <code>MB_CUR_MAX</code> macro.</p>								
ERRORS	<p>The <code>mblen()</code> function may fail if:</p> <p><code>EILSEQ</code> Invalid character sequence is detected.</p>								
USAGE	<p>The <code>mblen()</code> function can be used safely in multithreaded applications, as long as <code>setlocale(3C)</code> is not being called to change the locale.</p>								
ATTRIBUTES	<p>See <code>attributes(5)</code> 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>CSI</td> <td>Enabled</td> </tr> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	Standard	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
CSI	Enabled								
Interface Stability	Standard								
MT-Level	MT-Safe with exceptions								
SEE ALSO	<p><code>mbstowcs(3C)</code>, <code>mbtowc(3C)</code>, <code>setlocale(3C)</code>, <code>wcstombs(3C)</code>, <code>wctomb(3C)</code>, <code>attributes(5)</code>, <code>standards(5)</code></p>								

mbrlen(3C)

NAME	mbrlen – get number of bytes in a character (restartable)								
SYNOPSIS	<pre>#include <wchar.h> size_t mbrlen(const char *restrict s, size_t n, mbstate_t *restrict ps);</pre>								
DESCRIPTION	<p>If <i>s</i> is not a null pointer, <code>mbrlen()</code> determines the number of bytes constituting the character pointed to by <i>s</i>. It is equivalent to:</p> <pre>mbstate_t internal; mbrtowc(NULL, s, n, ps != NULL ? ps : &internal);</pre> <p>If <i>ps</i> is a null pointer, the <code>mbrlen()</code> function uses its own internal <code>mbstate_t</code> object, which is initialized at program startup to the initial conversion state. Otherwise, the <code>mbstate_t</code> object pointed to by <i>ps</i> is used to completely describe the current conversion state of the associated character sequence. Solaris will behave as if no function defined in the Solaris Reference Manual calls <code>mbrlen()</code>.</p> <p>The behavior of this function is affected by the <code>LC_CTYPE</code> category of the current locale. See <code>environ(5)</code>.</p>								
RETURN VALUES	<p>The <code>mbrlen()</code> function returns the first of the following that applies:</p> <table><tr><td>0</td><td>If the next <i>n</i> or fewer bytes complete the character that corresponds to the null wide-character.</td></tr><tr><td>positive</td><td>If the next <i>n</i> or fewer bytes complete a valid character; the value returned is the number of bytes that complete the character.</td></tr><tr><td>(size_t)-2</td><td>If the next <i>n</i> bytes contribute to an incomplete but potentially valid character, and all <i>n</i> bytes have been processed. When <i>n</i> has at least the value of the <code>MB_CUR_MAX</code> macro, this case can only occur if <i>s</i> points at a sequence of redundant shift sequences (for implementations with state-dependent encodings).</td></tr><tr><td>(size_t)-1</td><td>If an encoding error occurs, in which case the next <i>n</i> or fewer bytes do not contribute to a complete and valid character. In this case, <code>EILSEQ</code> is stored in <code>errno</code> and the conversion state is undefined.</td></tr></table>	0	If the next <i>n</i> or fewer bytes complete the character that corresponds to the null wide-character.	positive	If the next <i>n</i> or fewer bytes complete a valid character; the value returned is the number of bytes that complete the character.	(size_t)-2	If the next <i>n</i> bytes contribute to an incomplete but potentially valid character, and all <i>n</i> bytes have been processed. When <i>n</i> has at least the value of the <code>MB_CUR_MAX</code> macro, this case can only occur if <i>s</i> points at a sequence of redundant shift sequences (for implementations with state-dependent encodings).	(size_t)-1	If an encoding error occurs, in which case the next <i>n</i> or fewer bytes do not contribute to a complete and valid character. In this case, <code>EILSEQ</code> is stored in <code>errno</code> and the conversion state is undefined.
0	If the next <i>n</i> or fewer bytes complete the character that corresponds to the null wide-character.								
positive	If the next <i>n</i> or fewer bytes complete a valid character; the value returned is the number of bytes that complete the character.								
(size_t)-2	If the next <i>n</i> bytes contribute to an incomplete but potentially valid character, and all <i>n</i> bytes have been processed. When <i>n</i> has at least the value of the <code>MB_CUR_MAX</code> macro, this case can only occur if <i>s</i> points at a sequence of redundant shift sequences (for implementations with state-dependent encodings).								
(size_t)-1	If an encoding error occurs, in which case the next <i>n</i> or fewer bytes do not contribute to a complete and valid character. In this case, <code>EILSEQ</code> is stored in <code>errno</code> and the conversion state is undefined.								
ERRORS	<p>The <code>mbrlen()</code> function may fail if:</p> <table><tr><td><code>EINVAL</code></td><td>The <i>ps</i> argument points to an object that contains an invalid conversion state.</td></tr><tr><td><code>EILSEQ</code></td><td>Invalid character sequence is detected.</td></tr></table>	<code>EINVAL</code>	The <i>ps</i> argument points to an object that contains an invalid conversion state.	<code>EILSEQ</code>	Invalid character sequence is detected.				
<code>EINVAL</code>	The <i>ps</i> argument points to an object that contains an invalid conversion state.								
<code>EILSEQ</code>	Invalid character sequence is detected.								
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:								

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	See NOTES below

SEE ALSO [mbrtowc\(3C\)](#), [mbsinit\(3C\)](#), [setlocale\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#), [standards\(5\)](#)

NOTES If *ps* is not a null pointer, `mbrlen()` uses the `mbstate_t` object pointed to by *ps* and the function can be used safely in multithreaded applications, as long as [setlocale\(3C\)](#) is not being called to change the locale. If *ps* is a null pointer, `mbrlen()` uses its internal `mbstate_t` object and the function is Unsafe in multithreaded applications.

mbrtowc(3C)

NAME	mbrtowc – convert a character to a wide-character code (restartable)								
SYNOPSIS	<pre>#include <wchar.h> size_t mbrtowc(wchar_t *restrict pwc, const char *restrict s, size_t n, mbstate_t *restrict ps);</pre>								
DESCRIPTION	<p>If <i>s</i> is a null pointer, the <code>mbrtowc()</code> function is equivalent to the call:</p> <pre>mbrtowc(NULL, '', 1, ps)</pre> <p>In this case, the values of the arguments <i>pwc</i> and <i>n</i> are ignored.</p> <p>If <i>s</i> is not a null pointer, the <code>mbrtowc()</code> function inspects at most <i>n</i> bytes beginning at the byte pointed to by <i>s</i> to determine the number of bytes needed to complete the next character (including any shift sequences). If the function determines that the next character is completed, it determines the value of the corresponding wide-character and then, if <i>pwc</i> is not a null pointer, stores that value in the object pointed to by <i>pwc</i>. If the corresponding wide-character is the null wide-character, the resulting state described is the initial conversion state.</p> <p>If <i>ps</i> is a null pointer, the <code>mbrtowc()</code> function uses its own internal <code>mbstate_t</code> object, which is initialized at program startup to the initial conversion state. Otherwise, the <code>mbstate_t</code> object pointed to by <i>ps</i> is used to completely describe the current conversion state of the associated character sequence. Solaris will behave as if no function defined in the Solaris Reference Manual calls <code>mbrtowc()</code>.</p> <p>The behavior of this function is affected by the <code>LC_CTYPE</code> category of the current locale. See <code>environ(5)</code>.</p>								
RETURN VALUES	<p>The <code>mbrtowc()</code> function returns the first of the following that applies:</p> <table><tr><td>0</td><td>If the next <i>n</i> or fewer bytes complete the character that corresponds to the null wide-character (which is the value stored).</td></tr><tr><td>positive</td><td>If the next <i>n</i> or fewer bytes complete a valid character (which is the value stored); the value returned is the number of bytes that complete the character.</td></tr><tr><td>(size_t)-2</td><td>If the next <i>n</i> bytes contribute to an incomplete but potentially valid character, and all <i>n</i> bytes have been processed (no value is stored). When <i>n</i> has at least the value of the <code>MB_CUR_MAX</code> macro, this case can only occur if <i>s</i> points at a sequence of redundant shift sequences (for implementations with state-dependent encodings).</td></tr><tr><td>(size_t)-1</td><td>If an encoding error occurs, in which case the next <i>n</i> or fewer bytes do not contribute to a complete and valid character (no value is stored). In this case, <code>EILSEQ</code> is stored in <code>errno</code> and the conversion state is undefined.</td></tr></table>	0	If the next <i>n</i> or fewer bytes complete the character that corresponds to the null wide-character (which is the value stored).	positive	If the next <i>n</i> or fewer bytes complete a valid character (which is the value stored); the value returned is the number of bytes that complete the character.	(size_t)-2	If the next <i>n</i> bytes contribute to an incomplete but potentially valid character, and all <i>n</i> bytes have been processed (no value is stored). When <i>n</i> has at least the value of the <code>MB_CUR_MAX</code> macro, this case can only occur if <i>s</i> points at a sequence of redundant shift sequences (for implementations with state-dependent encodings).	(size_t)-1	If an encoding error occurs, in which case the next <i>n</i> or fewer bytes do not contribute to a complete and valid character (no value is stored). In this case, <code>EILSEQ</code> is stored in <code>errno</code> and the conversion state is undefined.
0	If the next <i>n</i> or fewer bytes complete the character that corresponds to the null wide-character (which is the value stored).								
positive	If the next <i>n</i> or fewer bytes complete a valid character (which is the value stored); the value returned is the number of bytes that complete the character.								
(size_t)-2	If the next <i>n</i> bytes contribute to an incomplete but potentially valid character, and all <i>n</i> bytes have been processed (no value is stored). When <i>n</i> has at least the value of the <code>MB_CUR_MAX</code> macro, this case can only occur if <i>s</i> points at a sequence of redundant shift sequences (for implementations with state-dependent encodings).								
(size_t)-1	If an encoding error occurs, in which case the next <i>n</i> or fewer bytes do not contribute to a complete and valid character (no value is stored). In this case, <code>EILSEQ</code> is stored in <code>errno</code> and the conversion state is undefined.								
ERRORS	The <code>mbrtowc()</code> function may fail if:								

EINVAL The *ps* argument points to an object that contains an invalid conversion state.

EILSEQ Invalid character sequence is detected.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	See NOTES below

SEE ALSO `mbsinit(3C)`, `setlocale(3C)`, `attributes(5)`, `environ(5)`, `standards(5)`

NOTES If *ps* is not a null pointer, `mbrtowc()` uses the `mbstate_t` object pointed to by *ps* and the function can be used safely in multithreaded applications, as long as `setlocale(3C)` is not being called to change the locale. If *ps* is a null pointer, `mbrtowc()` uses its internal `mbstate_t` object and the function is Unsafe in multithreaded applications.

mbsinit(3C)

NAME	mbsinit – determine conversion object status						
SYNOPSIS	<pre>#include <wchar.h> int mbsinit(const mbstate_t *ps);</pre>						
DESCRIPTION	If <i>ps</i> is not a null pointer, the <code>mbsinit()</code> function determines whether the object pointed to by <i>ps</i> describes an initial conversion state.						
RETURN VALUES	The <code>mbsinit()</code> function returns non-zero if <i>ps</i> is a null pointer, or if the pointed-to object describes an initial conversion state; otherwise, it returns 0. If an <code>mbstate_t</code> object is altered by any of the functions described as "restartable", and is then used with a different character sequence, or in the other conversion direction, or with a different <code>LC_CTYPE</code> category setting than on earlier function calls, the behavior is undefined. See <code>environ(5)</code> .						
ERRORS	No errors are defined.						
USAGE	The <code>mbstate_t</code> object is used to describe the current conversion state from a particular character sequence to a wide-character sequence (or vice versa) under the rules of a particular setting of the <code>LC_CTYPE</code> category of the current locale. The initial conversion state corresponds, for a conversion in either direction, to the beginning of a new character sequence in the initial shift state. A zero-valued <code>mbstate_t</code> object is at least one way to describe an initial conversion state. A zero-valued <code>mbstate_t</code> object can be used to initiate conversion involving any character sequence, in any <code>LC_CTYPE</code> category setting.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes: <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe with exceptions</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe with exceptions						
SEE ALSO	<code>mbrlen(3C)</code> , <code>mbrtowc(3C)</code> , <code>mbsrtowcs(3C)</code> , <code>setlocale(3C)</code> , <code>wcrtomb(3C)</code> , <code>wcsrtombs(3C)</code> , <code>attributes(5)</code> , <code>environ(5)</code> , <code>standards(5)</code>						
NOTES	The <code>mbsinit()</code> function can be used safely in multithreaded applications, as long as <code>setlocale(3C)</code> is not being called to change the locale.						

NAME	mbsrtowcs – convert a character string to a wide-character string (restartable)				
SYNOPSIS	<pre>#include <wchar.h> size_t mbsrtowcs(wchar_t *restrict dst, const char **restrict src, size_t len, mbstate_t *restrict ps);</pre>				
DESCRIPTION	<p>The <code>mbsrtowcs()</code> function converts a sequence of characters, beginning in the conversion state described by the object pointed to by <code>ps</code>, from the array indirectly pointed to by <code>src</code> into a sequence of corresponding wide-characters. If <code>dst</code> is not a null pointer, the converted characters are stored into the array pointed to by <code>dst</code>. Conversion continues up to and including a terminating null character, which is also stored. Conversion stops early in either of the following cases:</p> <ul style="list-style-type: none"> ■ When a sequence of bytes is encountered that does not form a valid character. ■ When <code>len</code> codes have been stored into the array pointed to by <code>dst</code> (and <code>dst</code> is not a null pointer). <p>Each conversion takes place as if by a call to the <code>mbrtowc()</code> function.</p> <p>If <code>dst</code> is not a null pointer, the pointer object pointed to by <code>src</code> is assigned either a null pointer (if conversion stopped due to reaching a terminating null character) or the address just past the last character converted (if any). If conversion stopped due to reaching a terminating null character, and if <code>dst</code> is not a null pointer, the resulting state described is the initial conversion state.</p> <p>If <code>ps</code> is a null pointer, the <code>mbsrtowcs()</code> function uses its own internal <code>mbstate_t</code> object, which is initialized at program startup to the initial conversion state. Otherwise, the <code>mbstate_t</code> object pointed to by <code>ps</code> is used to completely describe the current conversion state of the associated character sequence. Solaris will behave as if no function defined in the Solaris Reference Manual calls <code>mbsrtowcs()</code>.</p> <p>The behavior of this function is affected by the <code>LC_CTYPE</code> category of the current locale. See <code>environ(5)</code>.</p>				
RETURN VALUES	<p>If the input conversion encounters a sequence of bytes that do not form a valid character, an encoding error occurs. In this case, the <code>mbsrtowcs()</code> function stores the value of the macro <code>EILSEQ</code> in <code>errno</code> and returns <code>(size_t)-1</code>; the conversion state is undefined. Otherwise, it returns the number of characters successfully converted, not including the terminating null (if any).</p>				
ERRORS	<p>The <code>mbsrtowcs()</code> function may fail if:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;"><code>EINVAL</code></td> <td>The <code>ps</code> argument points to an object that contains an invalid conversion state.</td> </tr> <tr> <td><code>EILSEQ</code></td> <td>Invalid character sequence is detected.</td> </tr> </table>	<code>EINVAL</code>	The <code>ps</code> argument points to an object that contains an invalid conversion state.	<code>EILSEQ</code>	Invalid character sequence is detected.
<code>EINVAL</code>	The <code>ps</code> argument points to an object that contains an invalid conversion state.				
<code>EILSEQ</code>	Invalid character sequence is detected.				

mbsrtowcs(3C)

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	See NOTES below

SEE ALSO [mbrtowc\(3C\)](#), [mbsinit\(3C\)](#), [setlocale\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#), [standards\(5\)](#)

NOTES If *ps* is not a null pointer, `mbsrtowcs()` uses the `mbstate_t` object pointed to by *ps* and the function can be used safely in multithreaded applications, as long as [setlocale\(3C\)](#) is not being called to change the locale. If *ps* is a null pointer, `mbsrtowcs()` uses its internal `mbstate_t` object and the function is Unsafe in multithreaded applications.

NAME	mbstowcs – convert a character string to a wide-character string								
SYNOPSIS	<pre>#include <stdlib.h> size_t mbstowcs(wchar_t *restrict pwcs, const char *restrict s, size_t n);</pre>								
DESCRIPTION	<p>The <code>mbstowcs()</code> function converts a sequence of characters from the array pointed to by <code>s</code> into a sequence of corresponding wide-character codes and stores not more than <code>n</code> wide-character codes into the array pointed to by <code>pwcs</code>. No characters that follow a null byte (which is converted into a wide-character code with value 0) will be examined or converted. Each character is converted as if by a call to <code>mbtowc(3C)</code>.</p> <p>No more than <code>n</code> elements will be modified in the array pointed to by <code>pwcs</code>. If copying takes place between objects that overlap, the behavior is undefined.</p> <p>The behavior of this function is affected by the <code>LC_CTYPE</code> category of the current locale. If <code>pwcs</code> is a null pointer, <code>mbstowcs()</code> returns the length required to convert the entire array regardless of the value of <code>n</code>, but no values are stored.</p>								
RETURN VALUES	<p>If an invalid character is encountered, <code>mbstowcs()</code> returns <code>(size_t)-1</code> and may set <code>errno</code> to indicate the error. Otherwise, <code>mbstowcs()</code> returns the number of the array elements modified (or required if <code>pwcs</code> is <code>NULL</code>), not including a terminating 0 code, if any. The array will not be zero-terminated if the value returned is <code>n</code>.</p>								
ERRORS	<p>The <code>mbstowcs()</code> function may fail if:</p> <p><code>EILSEQ</code> Invalid byte sequence is detected.</p>								
ATTRIBUTES	<p>See <code>attributes(5)</code> 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>CSI</td> <td>Enabled</td> </tr> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
CSI	Enabled								
Interface Stability	Standard								
MT-Level	MT-Safe								
SEE ALSO	<p><code>mblen(3C)</code>, <code>mbtowc(3C)</code>, <code>setlocale(3C)</code>, <code>wcstombs(3C)</code>, <code>wctomb(3C)</code>, <code>attributes(5)</code>, <code>standards(5)</code></p>								

mbtowc(3C)

NAME	mbtowc – convert a character to a wide-character code								
SYNOPSIS	<pre>#include <stdlib.h> int mbtowc(wchar_t *restrict pwc, const char *restrict s, size_t n);</pre>								
DESCRIPTION	<p>If <i>s</i> is not a null pointer, <code>mbtowc()</code> determines the number of the bytes that constitute the character pointed to by <i>s</i>. It then determines the wide-character code for the value of type <code>wchar_t</code> that corresponds to that character. (The value of the wide-character code corresponding to the null byte is 0.) If the character is valid and <i>pwc</i> is not a null pointer, <code>mbtowc()</code> stores the wide-character code in the object pointed to by <i>pwc</i>.</p> <p>A call with <i>s</i> as a null pointer causes this function to return 0. The behavior of this function is affected by the <code>LC_CTYPE</code> category of the current locale. At most <i>n</i> bytes of the array pointed to by <i>s</i> will be examined.</p>								
RETURN VALUES	<p>If <i>s</i> is a null pointer, <code>mbtowc()</code> returns 0. If <i>s</i> is not a null pointer, <code>mbtowc()</code> returns 0 (if <i>s</i> points to the null byte), the number of bytes that constitute the converted character (if the next <i>n</i> or fewer bytes form a valid character), or -1 and may set <code>errno</code> to indicate the error (if they do not form a valid character).</p> <p>In no case will the value returned be greater than <i>n</i> or the value of the <code>MB_CUR_MAX</code> macro.</p>								
ERRORS	<p>The <code>mbtowc()</code> function may fail if:</p> <table><tr><td><code>EILSEQ</code></td><td>Invalid character sequence is detected.</td></tr></table>	<code>EILSEQ</code>	Invalid character sequence is detected.						
<code>EILSEQ</code>	Invalid character sequence is detected.								
USAGE	The <code>mbtowc()</code> function can be used safely in multithreaded applications, as long as <code>setlocale(3C)</code> is not being called to change the locale.								
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:								
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>CSI</td><td>Enabled</td></tr><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe with exceptions</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	Standard	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
CSI	Enabled								
Interface Stability	Standard								
MT-Level	MT-Safe with exceptions								
SEE ALSO	<code>mblen(3C)</code> , <code>mbstowcs(3C)</code> , <code>setlocale(3C)</code> , <code>wcstombs(3C)</code> , <code>wctomb(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>								

NAME	mctl – memory management control																				
SYNOPSIS	<pre> /usr/ucb/cc[flag ...] file ... #include <sys/types.h> #include <sys/mman.h> int mctl(addr, len, function, arg); caddr_t addr; size_t len; int function; int arg; </pre>																				
DESCRIPTION	<p>mctl() applies a variety of control functions over pages identified by the mappings established for the address range [addr, addr + len). The function to be performed is identified by the argument <i>function</i>. Valid functions are defined in mman.h as follows:</p> <table border="0"> <tr> <td style="vertical-align: top; padding-right: 20px;">MC_LOCK</td> <td>Lock the pages in the range in memory. This function is used to support mlock(). See mlock(3C) for semantics and usage. <i>arg</i> is ignored.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">MC_LOCKAS</td> <td>Lock the pages in the address space in memory. This function is used to support mlockall(). See mlockall(3C) for semantics and usage. <i>addr</i> and <i>len</i> are ignored. <i>arg</i> is an integer built from the flags: <table border="0" style="margin-left: 20px;"> <tr> <td>MCL_CURRENT</td> <td>Lock current mappings</td> </tr> <tr> <td>MCL_FUTURE</td> <td>Lock future mappings</td> </tr> </table> </td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">MC_SYNC</td> <td>Synchronize the pages in the range with their backing storage. Optionally invalidate cache copies. This function is used to support msync(). See msync(3C) for semantics and usage. <i>arg</i> is used to represent the <i>flags</i> argument to msync(). It is constructed from an OR of the following values: <table border="0" style="margin-left: 20px;"> <tr> <td>MS_SYNC</td> <td>Synchronized write</td> </tr> <tr> <td>MS_ASYNC</td> <td>Return immediately</td> </tr> <tr> <td>MS_INVALIDATE</td> <td>Invalidate mappings</td> </tr> </table> <p>MS_ASYNC returns after all I/O operations are scheduled. MS_SYNC does not return until all I/O operations are complete. Specify exactly one of MS_ASYNC or MS_SYNC. MS_INVALIDATE invalidates all cached copies of data from memory, requiring them to be re-obtained from the object's permanent storage location upon the next reference.</p> </td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">MC_UNLOCK</td> <td>Unlock the pages in the range. This function is used to support munlock(). <i>arg</i> is ignored.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">MC_UNLOCKAS</td> <td>Remove address space memory lock, and locks on all current mappings. This function is used to support munlockall(). <i>addr</i> and <i>len</i> must have the value 0. <i>arg</i> is ignored.</td> </tr> </table>	MC_LOCK	Lock the pages in the range in memory. This function is used to support mlock(). See mlock(3C) for semantics and usage. <i>arg</i> is ignored.	MC_LOCKAS	Lock the pages in the address space in memory. This function is used to support mlockall(). See mlockall(3C) for semantics and usage. <i>addr</i> and <i>len</i> are ignored. <i>arg</i> is an integer built from the flags: <table border="0" style="margin-left: 20px;"> <tr> <td>MCL_CURRENT</td> <td>Lock current mappings</td> </tr> <tr> <td>MCL_FUTURE</td> <td>Lock future mappings</td> </tr> </table>	MCL_CURRENT	Lock current mappings	MCL_FUTURE	Lock future mappings	MC_SYNC	Synchronize the pages in the range with their backing storage. Optionally invalidate cache copies. This function is used to support msync(). See msync(3C) for semantics and usage. <i>arg</i> is used to represent the <i>flags</i> argument to msync(). It is constructed from an OR of the following values: <table border="0" style="margin-left: 20px;"> <tr> <td>MS_SYNC</td> <td>Synchronized write</td> </tr> <tr> <td>MS_ASYNC</td> <td>Return immediately</td> </tr> <tr> <td>MS_INVALIDATE</td> <td>Invalidate mappings</td> </tr> </table> <p>MS_ASYNC returns after all I/O operations are scheduled. MS_SYNC does not return until all I/O operations are complete. Specify exactly one of MS_ASYNC or MS_SYNC. MS_INVALIDATE invalidates all cached copies of data from memory, requiring them to be re-obtained from the object's permanent storage location upon the next reference.</p>	MS_SYNC	Synchronized write	MS_ASYNC	Return immediately	MS_INVALIDATE	Invalidate mappings	MC_UNLOCK	Unlock the pages in the range. This function is used to support munlock(). <i>arg</i> is ignored.	MC_UNLOCKAS	Remove address space memory lock, and locks on all current mappings. This function is used to support munlockall(). <i>addr</i> and <i>len</i> must have the value 0. <i>arg</i> is ignored.
MC_LOCK	Lock the pages in the range in memory. This function is used to support mlock(). See mlock(3C) for semantics and usage. <i>arg</i> is ignored.																				
MC_LOCKAS	Lock the pages in the address space in memory. This function is used to support mlockall(). See mlockall(3C) for semantics and usage. <i>addr</i> and <i>len</i> are ignored. <i>arg</i> is an integer built from the flags: <table border="0" style="margin-left: 20px;"> <tr> <td>MCL_CURRENT</td> <td>Lock current mappings</td> </tr> <tr> <td>MCL_FUTURE</td> <td>Lock future mappings</td> </tr> </table>	MCL_CURRENT	Lock current mappings	MCL_FUTURE	Lock future mappings																
MCL_CURRENT	Lock current mappings																				
MCL_FUTURE	Lock future mappings																				
MC_SYNC	Synchronize the pages in the range with their backing storage. Optionally invalidate cache copies. This function is used to support msync(). See msync(3C) for semantics and usage. <i>arg</i> is used to represent the <i>flags</i> argument to msync(). It is constructed from an OR of the following values: <table border="0" style="margin-left: 20px;"> <tr> <td>MS_SYNC</td> <td>Synchronized write</td> </tr> <tr> <td>MS_ASYNC</td> <td>Return immediately</td> </tr> <tr> <td>MS_INVALIDATE</td> <td>Invalidate mappings</td> </tr> </table> <p>MS_ASYNC returns after all I/O operations are scheduled. MS_SYNC does not return until all I/O operations are complete. Specify exactly one of MS_ASYNC or MS_SYNC. MS_INVALIDATE invalidates all cached copies of data from memory, requiring them to be re-obtained from the object's permanent storage location upon the next reference.</p>	MS_SYNC	Synchronized write	MS_ASYNC	Return immediately	MS_INVALIDATE	Invalidate mappings														
MS_SYNC	Synchronized write																				
MS_ASYNC	Return immediately																				
MS_INVALIDATE	Invalidate mappings																				
MC_UNLOCK	Unlock the pages in the range. This function is used to support munlock(). <i>arg</i> is ignored.																				
MC_UNLOCKAS	Remove address space memory lock, and locks on all current mappings. This function is used to support munlockall(). <i>addr</i> and <i>len</i> must have the value 0. <i>arg</i> is ignored.																				

mctl(3UCB)

RETURN VALUES

mctl() returns 0 on success, -1 on failure.

ERRORS

mctl() fails if:

- EAGAIN Some or all of the memory identified by the operation could not be locked due to insufficient system resources.
- EBUSY MS_INVALIDATE was specified and one or more of the pages is locked in memory.
- EINVAL *addr* is not a multiple of the page size as returned by `getpagesize()`.
- EINVAL *addr* and/or *len* do not have the value 0 when MC_LOCKAS or MC_UNLOCKAS are specified.
- EINVAL *arg* is not valid for the function specified.
- EIO An I/O error occurred while reading from or writing to the file system.
- ENOMEM Addresses in the range [*addr*, *addr + len*) are invalid for the address space of a process, or specify one or more pages which are not mapped.
- EPERM The process's effective user ID is not super-user and one of MC_LOCK, MC_LOCKAS, MC_UNLOCK, or MC_UNLOCKAS was specified.

SEE ALSO

mmap(2), memcntl(2), [getpagesize\(3C\)](#), [mlock\(3C\)](#), [mlockall\(3C\)](#), [msync\(3C\)](#)

NOTES

Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.

NAME	memory, memccpy, memchr, memcmp, memcpy, memmove, memset – memory operations
SYNOPSIS	<pre>#include <string.h> void *memccpy(void *restrict s1, const void *restrict s2, int c, size_t n); void *memchr(const void *s, int c, size_t n); int memcmp(const void *s1, const void *s2, size_t n); void *memcpy(void *restrict s1, const void *restrict s2, size_t n); void *memmove(void *s1, const void *s2, size_t n); void *memset(void *s, int c, size_t n);</pre>
ISO C++	<pre>#include <string.h> const void *memchr(const void *s, int c, size_t n); #include <cstring> void *std::memchr(void *s, int c, size_t n);</pre>
DESCRIPTION	<p>These functions operate as efficiently as possible on memory areas (arrays of bytes bounded by a count, not terminated by a null character). They do not check for the overflow of any receiving memory area.</p> <p>The <code>memccpy()</code> function copies bytes from memory area <code>s2</code> into <code>s1</code>, stopping after the first occurrence of <code>c</code> (converted to an unsigned char) has been copied, or after <code>n</code> bytes have been copied, whichever comes first. It returns a pointer to the byte after the copy of <code>c</code> in <code>s1</code>, or a null pointer if <code>c</code> was not found in the first <code>n</code> bytes of <code>s2</code>.</p> <p>The <code>memchr()</code> function returns a pointer to the first occurrence of <code>c</code> (converted to an unsigned char) in the first <code>n</code> bytes (each interpreted as an unsigned char) of memory area <code>s</code>, or a null pointer if <code>c</code> does not occur.</p> <p>The <code>memcmp()</code> function compares its arguments, looking at the first <code>n</code> bytes (each interpreted as an unsigned char), and returns an integer less than, equal to, or greater than 0, according as <code>s1</code> is lexicographically less than, equal to, or greater than <code>s2</code> when taken to be unsigned characters.</p> <p>The <code>memcpy()</code> function copies <code>n</code> bytes from memory area <code>s2</code> to <code>s1</code>. It returns <code>s1</code>. If copying takes place between objects that overlap, the behavior is undefined.</p> <p>The <code>memmove()</code> function copies <code>n</code> bytes from memory area <code>s2</code> to memory area <code>s1</code>. Copying between objects that overlap will take place correctly. It returns <code>s1</code>.</p> <p>The <code>memset()</code> function sets the first <code>n</code> bytes in memory area <code>s</code> to the value of <code>c</code> (converted to an unsigned char). It returns <code>s</code>.</p>

memory(3C)

USAGE Using `memcpy()` might be faster than using `memmove()` if the application knows that the objects being copied do not overlap.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [string\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	mkfifo – make a FIFO special file																
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/stat.h> int mkfifo(const char *path, mode_t mode);</pre>																
DESCRIPTION	<p>The <code>mkfifo()</code> function creates a new FIFO special file named by the pathname pointed to by <i>path</i>. The file permission bits of the new FIFO are initialized from <i>mode</i>. The file permission bits of the <i>mode</i> argument are modified by the process's file creation mask (see <code>umask(2)</code>). Bits other than the file permission bits in <i>mode</i> are ignored.</p> <p>If <i>path</i> names a symbolic link, <code>mkfifo()</code> fails and sets <code>errno</code> to <code>EEXIST</code>.</p> <p>The FIFO's user ID is set to the process's effective user ID. The FIFO's group ID is set to the group ID of the parent directory or to the effective group ID of the process.</p> <p>The <code>mkfifo()</code> function calls <code>mknod(2)</code> to create the file.</p> <p>Upon successful completion, <code>mkfifo()</code> marks for update the <code>st_atime</code>, <code>st_ctime</code>, and <code>st_mtime</code> fields 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 and <code>errno</code> is set to indicate the error.																
ERRORS	<p>The <code>mkfifo()</code> function will fail if:</p> <table border="0"> <tr> <td style="vertical-align: top;">EACCES</td> <td>A component of the path prefix denies search permission, or write permission is denied on the parent directory of the FIFO to be created.</td> </tr> <tr> <td style="vertical-align: top;">EEXIST</td> <td>The named file already exists.</td> </tr> <tr> <td style="vertical-align: top;">ELOOP</td> <td>A loop exists in symbolic links encountered during resolution of the <i>path</i> argument.</td> </tr> <tr> <td style="vertical-align: top;">ENAMETOOLONG</td> <td>The length of the <i>path</i> argument exceeds <code>{PATH_MAX}</code> or a pathname component is longer than <code>{NAME_MAX}</code>.</td> </tr> <tr> <td style="vertical-align: top;">ENOENT</td> <td>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.</td> </tr> <tr> <td style="vertical-align: top;">ENOSPC</td> <td>The directory that would contain the new file cannot be extended or the file system is out of file-allocation resources.</td> </tr> <tr> <td style="vertical-align: top;">ENOTDIR</td> <td>A component of the path prefix is not a directory.</td> </tr> <tr> <td style="vertical-align: top;">EROFS</td> <td>The named file resides on a read-only file system.</td> </tr> </table> <p>The <code>mkfifo()</code> function may fail if:</p>	EACCES	A component of the path prefix denies search permission, or write permission is denied on the parent directory of the FIFO to be created.	EEXIST	The named file already exists.	ELOOP	A loop exists in symbolic links encountered during resolution of the <i>path</i> argument.	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> .	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.	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.	EROFS	The named file resides on a read-only file system.
EACCES	A component of the path prefix denies search permission, or write permission is denied on the parent directory of the FIFO to be created.																
EEXIST	The named file already exists.																
ELOOP	A loop exists in symbolic links encountered during resolution of the <i>path</i> argument.																
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> .																
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.																
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.																
EROFS	The named file resides on a read-only file system.																

mkfifo(3C)

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

ENAMETOOLONG The length of the *path* argument exceeds {PATH_MAX} or a pathname component is longer than {NAME_MAX}.

EXAMPLES

EXAMPLE 1 Create a FIFO File

The following example demonstrates how to create a FIFO file named /home/cnd/mod_done with read and write permissions for the owner and read permissions for the group and others.

```
#include sys/types.h>
#include sys/stat.h>
int status;
...
status = mkfifo("/home/cnd/mod_done", S_IWUSR | S_IRUSR |
               S_IRGRP | S_IROTH);
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO mkdir(1), chmod(2), exec(2), mknod(2), umask(2), stat.h(3HEAD), ufs(7FS), attributes(5), standards(5)

NAME	mkstemp – make a unique file name from a template and open the file				
SYNOPSIS	<pre>#include <stdlib.h> int mkstemp(char *<i>template</i>);</pre>				
DESCRIPTION	The <code>mkstemp()</code> function replaces the contents of the string pointed to by <i>template</i> by a unique file name, and returns a file descriptor for the file open for reading and writing. The function thus prevents any possible race condition between testing whether the file exists and opening it for use. The string in <i>template</i> should look like a file name with six trailing 'X's; <code>mkstemp()</code> replaces each 'X' with a character from the portable file name character set. The characters are chosen such that the resulting name does not duplicate the name of an existing file.				
RETURN VALUES	Upon successful completion, <code>mkstemp()</code> returns an open file descriptor. Otherwise <code>-1</code> is returned if no suitable file could be created.				
ERRORS	No errors are defined.				
USAGE	<p>It is possible to run out of letters.</p> <p>The <code>mkstemp()</code> function does not check to determine whether the file name part of <i>template</i> exceeds the maximum allowable file name length.</p> <p>The <code>tmpfile(3C)</code> function is preferred over this function.</p> <p>The <code>mkstemp()</code> function is frequently used to create a temporary file that will be removed by the application before the application terminates.</p> <p>The <code>mkstemp()</code> function has a transitional interface for 64-bit file offsets. See <code>lf64(5)</code>.</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Standard				
SEE ALSO	<code>getpid(2)</code> , <code>open(2)</code> , <code>tmpfile(3C)</code> , <code>tmpnam(3C)</code> , <code>attributes(5)</code> , <code>lf64(5)</code> , <code>standards(5)</code>				

mktemp(3C)

NAME	mktemp – make a unique file name from a template						
SYNOPSIS	<pre>#include <stdlib.h> char *mktemp(char *template);</pre>						
DESCRIPTION	The <code>mktemp()</code> function replaces the contents of the string pointed to by <i>template</i> with a unique file name, and returns <i>template</i> . The string in <i>template</i> should look like a file name with six trailing 'X's; <code>mktemp()</code> will replace the 'X's with a character string that can be used to create a unique file name. Only 26 unique file names per thread can be created for each unique <i>template</i> .						
RETURN VALUES	The <code>mktemp()</code> function returns the pointer <i>template</i> . If a unique name cannot be created, <i>template</i> points to a null string.						
ERRORS	No errors are defined.						
EXAMPLES	<p>EXAMPLE 1 Generate a filename.</p> <p>The following example replaces the contents of the "template" string with a 10-character filename beginning with the characters "file" and returns a pointer to the "template" string that contains the new filename.</p> <pre>#include <stdlib.h> ... char *template = "/tmp/fileXXXXXX"; char *ptr; ptr = mktemp(template);</pre>						
USAGE	Between the time a pathname is created and the file opened, it is possible for some other process to create a file with the same name. The <code>mkstemp(3C)</code> function avoids this problem and is preferred over this function.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	<code>mkstemp(3C)</code> , <code>tmpfile(3C)</code> , <code>tmpnam(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	mktime – converts a tm structure to a calendar time
SYNOPSIS	<pre>#include <time.h> time_t mktime(struct tm *timeptr);</pre>
DESCRIPTION	<p>The mktime() function converts the time represented by the tm structure pointed to by <i>timeptr</i> into a calendar time (the number of seconds since 00:00:00 UTC, January 1, 1970).</p> <p>The tm structure contains the following members:</p> <pre>int tm_sec; /* seconds after the minute [0, 60] */ int tm_min; /* minutes after the hour [0, 59] */ int tm_hour; /* hour since midnight [0, 23] */ int tm_mday; /* day of the month [1, 31] */ int tm_mon; /* months since January [0, 11] */ int tm_year; /* years since 1900 */ int tm_wday; /* days since Sunday [0, 6] */ int tm_yday; /* days since January 1 [0, 365] */ int tm_isdst; /* flag for daylight savings time */</pre> <p>In addition to computing the calendar time, mktime() normalizes the supplied tm structure. The original values of the tm_wday and tm_yday components of the structure are ignored, and the original values of the other components are not restricted to the ranges indicated in the definition of the structure. On successful completion, the values of the tm_wday and tm_yday components are set appropriately, and the other components are set to represent the specified calendar time, but with their values forced to be within the appropriate ranges. The final value of tm_mday is not set until tm_mon and tm_year are determined.</p> <p>The tm_year member must be for year 1901 or later. Calendar times before 20:45:52 UTC, December 13, 1901 or after 03:14:07 UTC, January 19, 2038 cannot be represented. Portable applications should not try to create dates before 00:00:00 UTC, January 1, 1970 or after 00:00:00 UTC, January 1, 2038.</p> <p>The original values of the components may be either greater than or less than the specified range. For example, a tm_hour of -1 means 1 hour before midnight, tm_mday of 0 means the day preceding the current month, and tm_mon of -2 means 2 months before January of tm_year.</p> <p>If tm_isdst is positive, the original values are assumed to be in the alternate timezone. If it turns out that the alternate timezone is not valid for the computed calendar time, then the components are adjusted to the main timezone. Likewise, if tm_isdst is zero, the original values are assumed to be in the main timezone and are converted to the alternate timezone if the main timezone is not valid. If tm_isdst is negative, mktime() attempts to determine whether the alternate timezone is in effect for the specified time.</p> <p>Local timezone information is used as if mktime() had called tzset(). See ctime(3C).</p>

mktime(3C)

RETURN VALUES	If the calendar time can be represented in an object of type <code>time_t</code> , <code>mktime()</code> returns the specified calendar time without changing <code>errno</code> . If the calendar time cannot be represented, the function returns the value <code>(time_t)-1</code> and sets <code>errno</code> to indicate the error.		
ERRORS	The <code>mktime()</code> function will fail if: <table><tr><td><code>EOVERFLOW</code></td><td>The date represented by the input <code>tm</code> struct cannot be represented in a <code>time_t</code>. Note that the <code>errno</code> setting may change if future revisions to the standards specify a different value.</td></tr></table>	<code>EOVERFLOW</code>	The date represented by the input <code>tm</code> struct cannot be represented in a <code>time_t</code> . Note that the <code>errno</code> setting may change if future revisions to the standards specify a different value.
<code>EOVERFLOW</code>	The date represented by the input <code>tm</code> struct cannot be represented in a <code>time_t</code> . Note that the <code>errno</code> setting may change if future revisions to the standards specify a different value.		
USAGE	<p>The <code>mktime()</code> function is MT-Safe in multithreaded applications, as long as no user-defined function directly modifies one of the following variables: <code>timezone</code>, <code>altzone</code>, <code>daylight</code>, and <code>tzname</code>. See ctime(3C).</p> <p>Note that <code>-1</code> can be a valid return value for the time that is one second before the Epoch. The user should clear <code>errno</code> before calling <code>mktime()</code>. If <code>mktime()</code> then returns <code>-1</code>, the user should check <code>errno</code> to determine whether or not an error actually occurred.</p> <p>The <code>mktime()</code> function assumes Gregorian dates. Times before the adoption of the Gregorian calendar will not match historical records.</p>		
EXAMPLES	<p>EXAMPLE 1 Sample code using <code>mktime()</code>.</p> <p>What day of the week is July 4, 2001?</p> <pre>#include <stdio.h> #include <time.h> static char *const wday[] = { "Sunday", "Monday", "Tuesday", "Wednesday", "Thursday", "Friday", "Saturday", "-unknown-" }; struct tm time_str; /* . . . */ time_str.tm_year = 2001 - 1900; time_str.tm_mon = 7 - 1; time_str.tm_mday = 4; time_str.tm_hour = 0; time_str.tm_min = 0; time_str.tm_sec = 1; time_str.tm_isdst = -1; if (mktime(&time_str) == -1) time_str.tm_wday = 7; printf("%s\n", wday[time_str.tm_wday]);</pre>		
BUGS	The <code>zoneinfo</code> <code>timezone</code> data files do not transition past Tue Jan 19 03:14:07 2038 UTC. Therefore for 64-bit applications using <code>zoneinfo</code> timezones, calculations beyond this date may not use the correct offset from standard time, and could return incorrect values. This affects the 64-bit version of <code>mktime()</code> .		

mktime(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

SEE ALSO `ctime(3C)`, `getenv(3C)`, `TIMEZONE(4)`, `attributes(5)`, `standards(5)`

mlock(3C)

NAME	mlock, munlock – lock or unlock pages in memory
SYNOPSIS	<pre>#include <sys/mman.h> int mlock(caddr_t <i>addr</i>, size_t <i>len</i>); int munlock(caddr_t <i>addr</i>, size_t <i>len</i>);</pre>
Standard conforming	<pre>#include <sys/mman.h> int mlock(const void * <i>addr</i>, size_t <i>len</i>); int munlock(const void * <i>addr</i>, size_t <i>len</i>);</pre>
DESCRIPTION	<p>The <code>mlock()</code> function uses the mappings established for the address range [<i>addr</i>, <i>addr</i> + <i>len</i>) to identify pages to be locked in memory. If the page identified by a mapping changes, such as occurs when a copy of a writable MAP_PRIVATE page is made upon the first store, the lock will be transferred to the newly copied private page.</p> <p>The <code>munlock()</code> function removes locks established with <code>mlock()</code>.</p> <p>A given page may be locked multiple times by executing an <code>mlock()</code> through different mappings. That is, if two different processes lock the same page, then the page will remain locked until both processes remove their locks. However, within a given mapping, page locks do not nest – multiple <code>mlock()</code> operations on the same address in the same process will all be removed with a single <code>munlock()</code>. Of course, a page locked in one process and mapped in another (or visible through a different mapping in the locking process) is still locked in memory. This fact can be used to create applications that do nothing other than lock important data in memory, thereby avoiding page I/O faults on references from other processes in the system.</p> <p>The contents of the locked pages will not be transferred to or from disk except when explicitly requested by one of the locking processes. This guarantee applies only to the mapped data, and not to any associated data structures (file descriptors and on-disk metadata, among others).</p> <p>If the mapping through which an <code>mlock()</code> has been performed is removed, an <code>munlock()</code> is implicitly performed. An <code>munlock()</code> is also performed implicitly when a page is deleted through file removal or truncation.</p> <p>Locks established with <code>mlock()</code> are not inherited by a child process after a <code>fork()</code> and are not nested.</p> <p>Attempts to <code>mlock()</code> more memory than a system-specific limit will fail.</p>
RETURN VALUES	Upon successful completion, the <code>mlock()</code> and <code>munlock()</code> functions return 0. Otherwise, no changes are made to any locks in the address space of the process, the functions return -1 and set <code>errno</code> to indicate the error.
ERRORS	The <code>mlock()</code> and <code>munlock()</code> functions will fail if: EINVAL The <i>addr</i> argument is not a multiple of the page size as returned by <code>sysconf(3C)</code> .

ENOMEM Addresses in the range [*addr*, *addr + len*) are invalid for the address space of a process, or specify one or more pages which are not mapped.

ENOSYS The system does not support this memory locking interface.

EPERM The {PRIV_PROC_LOCK_MEMORY} privilege is not asserted in the effective set of the calling process.

The `mlock()` function will fail if:

EAGAIN Some or all of the memory identified by the range [*addr*, *addr + len*) could not be locked because of insufficient system resources.

USAGE Because of the impact on system resources, the use of `mlock()` and `munlock()` is restricted to the superuser.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `fork(2)`, `memcntl(2)`, `mmap(2)`, `plock(3C)`, `mlockall(3C)`, `sysconf(3C)`, `attributes(5)`, `standards(5)`

mlockall(3C)

NAME	mlockall, munlockall – lock or unlock address space						
SYNOPSIS	<pre>#include <sys/mman.h> int mlockall(int flags); int munlockall(void);</pre>						
DESCRIPTION	<p>The <code>mlockall()</code> function locks in memory all pages mapped by an address space.</p> <p>The value of <i>flags</i> 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:</p> <pre> MCL_CURRENT Lock current mappings MCL_FUTURE Lock future mappings</pre> <p>If <code>MCL_FUTURE</code> is specified for <code>mlockall()</code>, mappings are locked as they are added to the address space (or replace existing mappings), provided sufficient memory is available. Locking in this manner is not persistent across the <code>exec</code> family of functions (see <code>exec(2)</code>).</p> <p>Mappings locked using <code>mlockall()</code> with any option may be explicitly unlocked with a <code>munlock()</code> call (see <code>mlock(3C)</code>).</p> <p>The <code>munlockall()</code> function removes address space locks and locks on mappings in the address space.</p> <p>All conditions and constraints on the use of locked memory that apply to <code>mlock(3C)</code> also apply to <code>mlockall()</code>.</p> <p>Locks established with <code>mlockall()</code> are not inherited by a child process after a <code>fork(2)</code> call, and are not nested.</p>						
RETURN VALUES	Upon successful completion, the <code>mlockall()</code> and <code>munlockall()</code> functions return 0. Otherwise, they return <code>-1</code> and set <code>errno</code> to indicate the error.						
ERRORS	The <code>mlockall()</code> and <code>munlockall()</code> functions will fail if:						
	<table><tr><td>EAGAIN</td><td>Some or all of the memory in the address space could not be locked due to sufficient resources. This error condition applies to <code>mlockall()</code> only.</td></tr><tr><td>EINVAL</td><td>The <i>flags</i> argument contains values other than <code>MCL_CURRENT</code> and <code>MCL_FUTURE</code>.</td></tr><tr><td>EPERM</td><td>The <code>{PRIV_PROC_LOCK_MEMORY}</code> privilege is not asserted in the effective set of the calling process.</td></tr></table>	EAGAIN	Some or all of the memory in the address space could not be locked due to sufficient resources. This error condition applies to <code>mlockall()</code> only.	EINVAL	The <i>flags</i> argument contains values other than <code>MCL_CURRENT</code> and <code>MCL_FUTURE</code> .	EPERM	The <code>{PRIV_PROC_LOCK_MEMORY}</code> privilege is not asserted in the effective set of the calling process.
EAGAIN	Some or all of the memory in the address space could not be locked due to sufficient resources. This error condition applies to <code>mlockall()</code> only.						
EINVAL	The <i>flags</i> argument contains values other than <code>MCL_CURRENT</code> and <code>MCL_FUTURE</code> .						
EPERM	The <code>{PRIV_PROC_LOCK_MEMORY}</code> privilege is not asserted in the effective set of the calling process.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						

mlockall(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [exec\(2\)](#), [fork\(2\)](#), [memcntl\(2\)](#), [mmap\(2\)](#), [plock\(3C\)](#), [mlock\(3C\)](#), [sysconf\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

monitor(3C)

NAME	monitor – prepare process execution profile
SYNOPSIS	<pre>#include <mon.h> void monitor(int (*lowpc()), int (*highpc()), WORD *buffer, size_t bufsize, size_t nfunc);</pre>
DESCRIPTION	<p>The <code>monitor()</code> function is an interface to the <code>profil(2)</code> function and is called automatically with default parameters by any program created by the <code>cc(1B)</code> utility with the <code>-p</code> option specified. Except to establish further control over profiling activity, it is not necessary to explicitly call <code>monitor()</code>.</p> <p>When used, <code>monitor()</code> is called at least at the beginning and the end of a program. The first call to <code>monitor()</code> initiates the recording of two different kinds of execution-profile information: execution-time distribution and function call count. Execution-time distribution data is generated by <code>profil()</code> and the function call counts are generated by code supplied to the object file (or files) by <code>cc(1B) -p</code>. Both types of information are collected as a program executes. The last call to <code>monitor()</code> writes this collected data to the output file <code>mon.out</code>.</p> <p>The name of the file written by <code>monitor()</code> is controlled by the environment variable <code>PROFDIR</code>. If <code>PROFDIR</code> does not exist, the file <code>mon.out</code> is created in the current directory. If <code>PROFDIR</code> exists but has no value, <code>monitor()</code> does no profiling and creates no output file. If <code>PROFDIR</code> is <code>dirname</code>, and <code>monitor()</code> is called automatically by compilation with <code>cc -p</code>, the file created is <code>dirname/pid.progname</code> where <i>progname</i> is the name of the program.</p> <p>The <i>lowpc</i> and <i>highpc</i> arguments are the beginning and ending addresses of the region to be profiled.</p> <p>The <i>buffer</i> argument is the address of a user-supplied array of <code>WORD</code> (defined in the header <code><mon.h></code>). The <i>buffer</i> argument is used by <code>monitor()</code> to store the histogram generated by <code>profil()</code> and the call counts.</p> <p>The <i>bufsize</i> argument identifies the number of array elements in <i>buffer</i>.</p> <p>The <i>nfunc</i> argument is the number of call count cells that have been reserved in <i>buffer</i>. Additional call count cells will be allocated automatically as they are needed.</p> <p>The <i>bufsize</i> argument should be computed using the following formula:</p> <pre>size_of_buffer = sizeof(struct hdr) + nfunc * sizeof(struct cnt) + ((highpc-lowpc)/BARSIZE) * sizeof(WORD) + sizeof(WORD) - 1 ; bufsize = (size_of_buffer / sizeof(WORD));</pre> <p>where:</p> <ul style="list-style-type: none">■ <i>lowpc</i>, <i>highpc</i>, <i>nfunc</i> are the same as the arguments to <code>monitor()</code>;■ <i>BARSIZE</i> is the number of program bytes that correspond to each histogram bar, or cell, of the <code>profil()</code> buffer;

- the `hdr` and `cnt` structures and the type `WORD` are defined in the header `<mon.h>`.

The default call to `monitor()` is as follows:

```
monitor (&eprol, &etext, wbuf, wbufsz, 600);
```

where:

- `eprol` is the beginning of the user's program when linked with `cc -p` (see [end\(3C\)](#));
- `etext` is the end of the user's program (see [end\(3C\)](#));
- `wbuf` is an array of `WORD` with `wbufsz` elements;
- `wbufsz` is computed using the `bufsize` formula shown above with `BARSIZE` of 8;
- 600 is the number of call count cells that have been reserved in `buffer`.

These parameter settings establish the computation of an execution-time distribution histogram that uses `profil()` for the entire program, initially reserves room for 600 call count cells in `buffer`, and provides for enough histogram cells to generate significant distribution-measurement results. For more information on the effects of `bufsize` on execution-distribution measurements, see [profil\(2\)](#).

EXAMPLES **EXAMPLE 1** Example to stop execution monitoring and write the results to a file.

To stop execution monitoring and write the results to a file, use the following:

```
monitor( (int (*)())0, (int (*)())0, (WORD *)0, 0, 0);
```

Use `prof` to examine the results.

USAGE Additional calls to `monitor()` after `main()` has been called and before `exit()` has been called will add to the function-call count capacity, but such calls will also replace and restart the `profil()` histogram computation.

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO [cc\(1B\)](#), [profil\(2\)](#), [end\(3C\)](#), [attributes\(5\)](#), [prof\(5\)](#)

msync(3C)

NAME	msync – synchronize memory with physical storage						
SYNOPSIS	<pre>#include <sys/mman.h> int msync(void *addr, size_t len, int flags);</pre>						
DESCRIPTION	<p>The <code>msync()</code> function writes all modified copies of pages over the range <code>[addr, addr + len)</code> to the underlying hardware, or invalidates any copies so that further references to the pages will be obtained by the system from their permanent storage locations. The permanent storage for a modified <code>MAP_SHARED</code> mapping is the file the page is mapped to; the permanent storage for a modified <code>MAP_PRIVATE</code> mapping is its swap area.</p> <p>The <code>flags</code> argument is a bit pattern built from the following values:</p> <table><tr><td><code>MS_ASYNC</code></td><td>perform asynchronous writes</td></tr><tr><td><code>MS_SYNC</code></td><td>perform synchronous writes</td></tr><tr><td><code>MS_INVALIDATE</code></td><td>invalidate mappings</td></tr></table> <p>If <code>flags</code> is <code>MS_ASYNC</code> or <code>MS_SYNC</code>, the function synchronizes the file contents to match the current contents of the memory region.</p> <ul style="list-style-type: none">■ All write references to the memory region made prior to the call are visible by subsequent read operations on the file.■ All writes to the same portion of the file prior to the call may or may not be visible by read references to the memory region.■ Unmodified pages in the specified range are not written to the underlying hardware. <p>If <code>flags</code> is <code>MS_ASYNC</code>, the function may return immediately once all write operations are scheduled; if <code>flags</code> is <code>MS_SYNC</code>, the function does not return until all write operations are completed.</p> <p>If <code>flags</code> is <code>MS_INVALIDATE</code>, the function synchronizes the contents of the memory region to match the current file contents.</p> <ul style="list-style-type: none">■ All writes to the mapped portion of the file made prior to the call are visible by subsequent read references to the mapped memory region.■ All write references prior to the call, by any process, to memory regions mapped to the same portion of the file using <code>MAP_SHARED</code>, are visible by read references to the region. <p>If <code>msync()</code> causes any write to the file, then the file's <code>st_ctime</code> and <code>st_mtime</code> fields are marked for update.</p>	<code>MS_ASYNC</code>	perform asynchronous writes	<code>MS_SYNC</code>	perform synchronous writes	<code>MS_INVALIDATE</code>	invalidate mappings
<code>MS_ASYNC</code>	perform asynchronous writes						
<code>MS_SYNC</code>	perform synchronous writes						
<code>MS_INVALIDATE</code>	invalidate mappings						
RETURN VALUES	Upon successful completion, <code>msync()</code> returns 0; otherwise, it returns -1 and sets <code>errno</code> to indicate the error.						
ERRORS	The <code>msync()</code> function will fail if:						

- EBUSY Some or all of the addresses in the range [*addr*, *addr + len*) are locked and MS_SYNC with the MS_INVALIDATE option is specified.
- EAGAIN Some or all pages in the range [*addr*, *addr + len*) are locked for I/O.
- EINVAL The *addr* argument is not a multiple of the page size as returned by [sysconf\(3C\)](#).

The *flags* argument is not some combination of MS_ASYNC and MS_INVALIDATE.
- EIO An I/O error occurred while reading from or writing to the file system.
- ENOMEM Addresses in the range [*addr*, *addr + len*) are outside the valid range for the address space of a process, or specify one or more pages that are not mapped.
- EPERM MS_INVALIDATE was specified and one or more of the pages is locked in memory.

USAGE The `msync()` function should be used by programs that require a memory object to be in a known state, for example in building transaction facilities.

Normal system activity can cause pages to be written to disk. Therefore, there are no guarantees that `msync()` is the only control over when pages are or are not written to disk.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `memcntl(2)`, `mmap(2)`, [sysconf\(3C\)](#), `attributes(5)`, `standards(5)`

mtmalloc(3MALLOC)

NAME	mtmalloc, mallocctl – MT hot memory allocator
SYNOPSIS	<pre>#include <mtmalloc.h> cc -o a.out -lthread -lmtmalloc void *malloc(size_t size); void free(void *ptr); void *memalign(size_t alignment, size_t size); void *realloc(void *ptr, size_t size); void *valloc(size_t size); void mallocctl(int cmd, long value);</pre>
DESCRIPTION	<p>The <code>malloc()</code> and <code>free()</code> functions provide a simple general-purpose memory allocation package that is suitable for use in high performance multithreaded applications. The suggested use of this library is in multithreaded applications; it can be used for single threaded applications, but there is no advantage in doing so. This library cannot be dynamically loaded with <code>dlopen()</code> during runtime because there must be only one manager of the process heap.</p> <p>The <code>malloc()</code> function returns a pointer to a block of at least <i>size</i> bytes suitably aligned for any use.</p> <p>The argument to <code>free()</code> is a pointer to a block previously allocated by <code>malloc()</code> or <code>realloc()</code>. After <code>free()</code> is performed this space is available for further allocation. If <i>ptr</i> is a null pointer, no action occurs.</p> <p>Undefined results will occur if the space assigned by <code>malloc()</code> is overrun or if a random number is handed to <code>free()</code>. A freed pointer that is passed to <code>free()</code> will send a SIGABRT signal to the calling process. This behavior is controlled by <code>mallocctl()</code>.</p> <p>The <code>memalign()</code> function allocates <i>size</i> bytes on a specified alignment boundary and returns a pointer to the allocated block. The value of the returned address is guaranteed to be an even multiple of <i>alignment</i>. Note that the value of <i>alignment</i> must be a power of two, and must be greater than or equal to the size of a word.</p> <p>The <code>realloc()</code> function changes the size of the block pointed to by <i>ptr</i> to <i>size</i> bytes and returns a pointer to the (possibly moved) block. The contents will be unchanged up to the lesser of the new and old sizes. If the new size of the block requires movement of the block, the space for the previous instantiation of the block is freed. If the new size is larger, the contents of the newly allocated portion of the block are unspecified. If <i>ptr</i> is NULL, <code>realloc()</code> behaves like <code>malloc()</code> for the specified size. If <i>size</i> is 0 and <i>ptr</i> is not a null pointer, the space pointed to is freed.</p> <p>The <code>valloc()</code> function has the same effect as <code>malloc()</code>, except that the allocated memory will be aligned to a multiple of the value returned by <code>sysconf(_SC_PAGESIZE)</code>.</p>

After possible pointer coercion, each allocation routine returns a pointer to a space that is suitably aligned for storage of any type of object.

The `malloc()`, `realloc()`, `memalign()`, and `valloc()` functions will fail if there is not enough available memory.

The `mallocctl()` function controls the behavior of the `malloc` library. The options fall into two general classes, debugging options and performance options.

MTDOUBLEFREE	Allows double <code>free</code> of a pointer. Setting <i>value</i> to 1 means yes and 0 means no. The default behavior of double <code>free</code> results in a core dump.
MTDEBUGPATTERN	Writes misaligned data into the buffer after <code>free()</code> . When the buffer is reallocated, the contents are verified to ensure that there was no access to the buffer after the <code>free</code> . If the buffer has been dirtied, a <code>SIGABRT</code> signal is delivered to the process. Setting <i>value</i> to 1 means yes and 0 means no. The default behavior is to <i>not</i> write misaligned data. The pattern used is <code>0xdeadbeef</code> . Use of this option results in a performance penalty.
MTINITBUFFER	Writes misaligned data into the newly allocated buffer. This option is useful for detecting some accesses before initialization. Setting <i>value</i> to 1 means yes and 0 means no. The default behavior is to <i>not</i> write misaligned data to the newly allocated buffer. The pattern used is <code>0xbaddcafe</code> . Use of this option results in a performance penalty.
MTCHUNKSIZE	This option changes the size of allocated memory when a pool has exhausted all available memory in the buffer. Increasing this value allocates more memory for the application. A substantial performance gain can occur because the library makes fewer calls to the OS for more memory. Acceptable number <i>values</i> are between 9 and 256; the default value is 9. This value is multiplied by 8192.

RETURN VALUES If there is no available memory, `malloc()`, `realloc()`, `memalign()`, and `valloc()` return a null pointer. When `realloc()` is called with *size* > 0 and returns `NULL`, the block pointed to by *ptr* is left intact. If *size*, *nelem*, or *elsize* is 0, either a null pointer or a unique pointer that can be passed to `free()` is returned.

If `malloc()` or `realloc()` returns unsuccessfully, `errno` will be set to indicate the error.

ERRORS The `malloc()` and `realloc()` functions will fail if:

ENOMEM	The physical limits of the system are exceeded by <i>size</i> bytes of memory which cannot be allocated.
--------	--

mtmalloc(3MALLOC)

EAGAIN There is not enough memory available to allocate *size* bytes of memory; but the application could try again later.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `brk(2)`, `getrlimit(2)`, `bsdmalloc(3MALLOC)`, `dlopen(3C)`, `malloc(3C)`, `malloc(3MALLOC)`, `mapmalloc(3MALLOC)`, `signal.h(3HEAD)`, `watchmalloc(3MALLOC)`, `attributes(5)`

WARNINGS Undefined results will occur if the size requested for a block of memory exceeds the maximum size of a process's heap. This information may be obtained using `getrlimit()`.

NOTES Comparative Features of `malloc(3C)`, `bsdmalloc(3MALLOC)`, `malloc(3MALLOC)`, and `mtmalloc`.

- The `bsdmalloc(3MALLOC)` routines afford better performance, but are space-inefficient.
- The `malloc(3MALLOC)` routines are space-efficient, but have slower performance.
- The standard, fully SCD-compliant `malloc` routines are a trade-off between performance and space-efficiency.
- The `mtmalloc` routines provide fast, concurrent `malloc()` implementation that is space-inefficient.

The `free()` function does not set `errno`.

NAME	mutex_init, mutex_destroy, mutex_lock, mutex_trylock, mutex_unlock – mutual exclusion locks				
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <thread.h> #include <synch.h> int mutex_init(mutex_t *mp, int type, void * arg); int mutex_lock(mutex_t *mp); int mutex_trylock(mutex_t *mp); int mutex_unlock(mutex_t *mp); int mutex_destroy(mutex_t *mp);</pre>				
DESCRIPTION	<p>Mutual exclusion locks (mutexes) prevent multiple threads from simultaneously executing critical sections of code which access shared data (that is, mutexes are used to serialize the execution of threads). All mutexes must be global. A successful call for a mutex lock by way of <code>mutex_lock()</code> will cause another thread that is also trying to lock the same mutex to block until the owner thread unlocks it by way of <code>mutex_unlock()</code>. Threads within the same process or within other processes can share mutexes.</p> <p>Mutexes can synchronize threads within the same process or in other processes. Mutexes can be used to synchronize threads between processes if the mutexes are allocated in writable memory and shared among the cooperating processes (see <code>mmap(2)</code>), and have been initialized for this task.</p>				
Initialize	<p>Mutexes are either intra-process or inter-process, depending upon the argument passed implicitly or explicitly to the initialization of that mutex. A statically allocated mutex does not need to be explicitly initialized; by default, a statically allocated mutex is initialized with all zeros and its scope is set to be within the calling process.</p> <p>For inter-process synchronization, a mutex needs to be allocated in memory shared between these processes. Since the memory for such a mutex must be allocated dynamically, the mutex needs to be explicitly initialized using <code>mutex_init()</code>.</p> <p>The <code>mutex_init()</code> function initializes the mutex referenced by <code>mp</code> with the type specified by <code>type</code>. Upon successful initialization the state of the mutex becomes initialized and unlocked. No current type uses <code>arg</code> although a future type may specify additional behavior parameters by way of <code>arg</code>. The <code>type</code> argument can be one of the following:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">USYNC_THREAD</td> <td>The mutex can synchronize threads only in this process. The <code>arg</code> argument is ignored.</td> </tr> <tr> <td>USYNC_PROCESS</td> <td>The mutex can synchronize threads in this process and other processes. The <code>arg</code> argument is ignored. The object initialized with this attribute must be allocated in memory shared between processes, either in System V shared memory (see <code>shmop(2)</code>) or in memory</td> </tr> </table>	USYNC_THREAD	The mutex can synchronize threads only in this process. The <code>arg</code> argument is ignored.	USYNC_PROCESS	The mutex can synchronize threads in this process and other processes. The <code>arg</code> argument is ignored. The object initialized with this attribute must be allocated in memory shared between processes, either in System V shared memory (see <code>shmop(2)</code>) or in memory
USYNC_THREAD	The mutex can synchronize threads only in this process. The <code>arg</code> argument is ignored.				
USYNC_PROCESS	The mutex can synchronize threads in this process and other processes. The <code>arg</code> argument is ignored. The object initialized with this attribute must be allocated in memory shared between processes, either in System V shared memory (see <code>shmop(2)</code>) or in memory				

mutex_init(3C)

USYNC_PROCESS_ROBUST	mapped to a file (see <code>mmap(2)</code>). If the object is not allocated in such shared memory, it will not be shared between processes.
	<p>The mutex can synchronize threads in this process and other processes robustly. At the time of process death, if the lock is held by the process, it is unlocked. The next owner of this mutex will acquire it with an error return of <code>EOWNERDEAD</code>. The application must always check the return value from <code>mutex_lock()</code> for a mutex of this type. The new owner of this mutex should then attempt to make the state protected by the mutex consistent, since this state could have been left inconsistent when the last owner died. If the new owner is able to make the state consistent, it should reinitialize the mutex by calling <code>mutex_init()</code> and then unlock the mutex. Only the new owner can make the mutex consistent. If another process then calls <code>mutex_init()</code>, the call will return <code>EBUSY</code>. If for any reason the new owner is not able to make the state consistent, it should not reinitialize the mutex but should simply unlock the mutex. All waiting processes will be awakened and all subsequent calls to <code>mutex_lock()</code> will fail in acquiring the mutex with an error value of <code>ENOTRECOVERABLE</code>. The mutex can be reused by uninitialized the mutex with <code>mutex_destroy()</code> and reinitializing it with <code>mutex_init()</code>. If the process that acquired the lock with <code>EOWNERDEAD</code> died, the next owner will acquire the lock with an error value of <code>EOWNERDEAD</code>. The <i>arg</i> argument is ignored. The object initialized with this attribute must be allocated in memory shared between processes, either in System V shared memory (see <code>shmop(2)</code>) or in memory mapped to a file (see <code>mmap(2)</code>) and memory must be zeroed before initialization. All the processes interested in the robust lock must call <code>mutex_init()</code> at least once to register robust mutex with the system and potentially initialize it. If the object is not allocated in such shared memory, it will not be shared between processes. If <code>mutex_init()</code> is called on a previously initialized mutex, <code>mutex_init()</code> will not reinitialize the mutex and will return the error value <code>EBUSY</code>.</p> <p>If <i>type</i> is either <code>USYNC_THREAD</code> or <code>USYNC_PROCESS</code>, it can be augmented by the bitwise-inclusive-OR of one or both of the following flags:</p>

LOCK_RECURSIVE	A thread attempting to relock this mutex without first unlocking it will succeed in locking the mutex.
LOCK_ERRORCHECK	Unless LOCK_RECURSIVE is also set, a thread attempting to relock this mutex without first unlocking it will return with an error rather than deadlocking itself. A thread attempting to unlock this mutex without first owning it will return with an error.

See [pthread_mutexattr_settype\(3C\)](#) for more information on recursive and error checking mutex types. The combination (LOCK_RECURSIVE | LOCK_ERRORCHECK) is the same as the POSIX PTHREAD_MUTEX_RECURSIVE type.

Initializing mutexes can also be accomplished by allocating in zeroed memory (default), in which case a *type* of USYNC_THREAD is assumed. In general, the following rules apply to mutex initialization:

- The same mutex must not be simultaneously initialized by multiple threads.
- A mutex lock must not be reinitialized while in use by other threads.

These rules do not apply to USYNC_PROCESS_ROBUST mutexes. See the description for USYNC_PROCESS_ROBUST above. If default mutex attributes are used, the macro DEFAULTMUTEX can be used to initialize mutexes that are statically allocated.

Default mutex initialization (intra-process):

```
mutex_t mp;
mutex_init(&mp, NULL, NULL);
```

or

```
mutex_init(&mp, USYNC_THREAD, NULL);
```

or

```
mutex_t mp = DEFAULTMUTEX;
```

or

```
mutex_t *mp;
mp = calloc(1, sizeof (mutex_t));
```

or

```
mutex_t *mp;
mp = malloc(sizeof (mutex_t));
memset(mp, 0, sizeof (mutex_t));
```

Customized mutex initialization (inter-process):

```
mutex_init(&mp, USYNC_PROCESS, NULL);
```

Customized mutex initialization (inter-process):

mutex_init(3C)

```
mutex_init(&mp, USYNC_PROCESS_ROBUST, NULL);
```

Statically allocated mutexes can also be initialized with macros specifying `LOCK_RECURSIVE` and/or `LOCK_ERRORCHECK`:

```
mutex_t mp = RECURSIVEMUTEX;  
    Same as (USYNC_THREAD | LOCK_RECURSIVE)
```

```
mutex_t mp = ERRORCHECKMUTEX;  
    Same as (USYNC_THREAD | LOCK_ERRORCHECK)
```

```
mutex_t mp = RECURSIVE_ERRORCHECKMUTEX;  
    Same as (USYNC_THREAD | LOCK_RECURSIVE | LOCK_ERRORCHECK)
```

Lock and Unlock

A critical section of code is enclosed by a the call to lock the mutex and the call to unlock the mutex to protect it from simultaneous access by multiple threads. Only one thread at a time may possess mutually exclusive access to the critical section of code that is enclosed by the mutex-locking call and the mutex-unlocking call, whether the mutex's scope is intra-process or inter-process. A thread calling to lock the mutex either gets exclusive access to the code starting from the successful locking until its call to unlock the mutex, or it waits until the mutex is unlocked by the thread that locked it.

Mutexes have ownership, unlike semaphores. Although any thread, within the scope of a mutex, can get an unlocked mutex and lock access to the same critical section of code, only the thread that locked a mutex should unlock it.

If a thread waiting for a mutex receives a signal, upon return from the signal handler, the thread resumes waiting for the mutex as if there was no interrupt. A mutex protects code, not data; therefore, strongly bind a mutex with the data by putting both within the same structure, or at least within the same procedure.

A call to `mutex_lock()` locks the mutex object referenced by *mp*. If the mutex is already locked, the calling thread blocks until the mutex is freed; this will return with the mutex object referenced by *mp* in the locked state with the calling thread as its owner. If the current owner of a mutex tries to relock the mutex, it will result in deadlock.

The `mutex_trylock()` function is the same as `mutex_lock()`, respectively, except that if the mutex object referenced by *mp* is locked (by any thread, including the current thread), the call returns immediately with an error.

The `mutex_unlock()` function are called by the owner of the mutex object referenced by *mp* to release it. The mutex must be locked and the calling thread must be the one that last locked the mutex (the owner). If there are threads blocked on the mutex object referenced by *mp* when `mutex_unlock()` is called, the *mp* is freed, and the scheduling policy will determine which thread gets the mutex. If the calling thread is not the owner of the lock, no error status is returned, and the behavior of the program is undefined.

Destroy	The <code>mutex_destroy()</code> function destroys the mutex object referenced by <code>mp</code> . The mutex object becomes uninitialized. The space used by the destroyed mutex variable is not freed. It needs to be explicitly reclaimed.
RETURN VALUES	If successful, these functions return 0. Otherwise, an error number is returned.
ERRORS	<p>These functions may fail if:</p> <p>EFAULT The <code>mp</code> argument points to an illegal address.</p> <p>The <code>mutex_init()</code> function will fail if:</p> <p>EINVAL The value specified by <code>type</code> is invalid.</p> <p>The <code>mutex_init()</code> function will fail for <code>USYNC_PROCESS_ROBUST</code> type mutex if:</p> <p>EBUSY The mutex pointed to by <code>mp</code> was already initialized. An attempt to reinitialize a mutex previously initialized, but not yet destroyed.</p> <p>The <code>mutex_trylock()</code> function will fail if:</p> <p>EBUSY The mutex pointed to by <code>mp</code> was already locked.</p> <p>The <code>mutex_lock()</code> and <code>mutex_trylock()</code> functions will fail for a <code>LOCK_RECURSIVE</code> mutex if:</p> <p>EAGAIN The mutex could not be acquired because the maximum number of recursive locks for the mutex has been reached.</p> <p>The <code>mutex_lock()</code> function will fail for a <code>LOCK_ERRORCHECK</code> and non-<code>LOCK_RECURSIVE</code> mutex if:</p> <p>EDEADLK The current thread already owns the mutex.</p> <p>The <code>mutex_unlock()</code> function will fail for a <code>LOCK_ERRORCHECK</code> mutex if:</p> <p>EPERM The current thread does not own the mutex.</p> <p>The <code>mutex_lock()</code> or <code>mutex_trylock()</code> functions will fail for <code>USYNC_PROCESS_ROBUST</code> type mutex if:</p> <p>EOWNERDEAD The last owner of this mutex died while holding the mutex. This mutex is now owned by the caller. The caller must now attempt to make the state protected by the mutex consistent. If it is able to clean up the state, then it should reinitialize the mutex by calling <code>mutex_init()</code> and unlock the mutex. Subsequent calls to <code>mutex_lock()</code> will behave normally, as before. If the caller is not able to clean up the state, the mutex should not be reinitialized but it should be unlocked. Subsequent calls to <code>mutex_lock()</code> will fail to acquire the mutex, returning with the error value</p>

mutex_init(3C)

	ENOTRECOVERABLE. If the owner who acquired the lock with EOWNERDEAD died, the next owner will acquire the lock with EOWNERDEAD.
ELOCKUNMAPPED	The last owner of this mutex unmaped the mutex while holding the mutex. This mutex is now owned by the caller. The caller must now attempt to make the state protected by the mutex consistent. If it is able to clean up the state, it should reinitialize and unlock the mutex. Subsequent calls to <code>mutex_lock()</code> will behave normally, as before. If the caller is not able to clean up the state, the mutex should not be reinitialized. Subsequent calls to <code>mutex_lock()</code> will fail to acquire the mutex and return the error value ENOTRECOVERABLE.
ENOTRECOVERABLE	The mutex trying to be acquired is protecting the state that has been left irrecoverable when the mutex's last owner could not make the state protected by the mutex consistent. The mutex has not been acquired. This condition can occur when the lock was previously acquired with EOWNERDEAD or ELOCKUNMAPPED and the owner was not able to clean up the state and unlocked the mutex without making the mutex consistent.

EXAMPLES

Single Gate

The following example uses one global mutex as a gate-keeper to permit each thread exclusive sequential access to the code within the user-defined function "change_global_data." This type of synchronization will protect the state of shared data, but it also prohibits parallelism.

```
/* cc thisfile.c -pthread */
#define _REENTRANT
#include <stdio.h>
#include <thread.h>
#define NUM_THREADS 12
void *change_global_data(void *); /* for thr_create() */
main(int argc, char * argv[] ) {
    int i=0;
    for (i=0; i< NUM_THREADS; i++) {
        thr_create(NULL, 0, change_global_data, NULL, 0, NULL);
    }
    while ((thr_join(NULL, NULL, NULL) == 0));
}

void * change_global_data(void *null){
    static mutex_t Global_mutex;
    static int Global_data = 0;
    mutex_lock(&Global_mutex);
    Global_data++;
}
```


**Multiple
Instruction Single
Data**

```

        sleep(1);
        printf("%d is global data\n",Global_data);
        mutex_unlock(&Global_mutex);
        return NULL;
    }

```

The previous example, the mutex, the code it owns, and the data it protects was enclosed in one function. The next example uses C++ features to accommodate many functions that use just one mutex to protect one data:

```

/* CC thisfile.c -lthread use C++ to compile*/

#define _REENTRANT
#include <stdlib.h>
#include <stdio.h>
#include <thread.h>
#include <errno.h>
#include <iostream.h>
#define NUM_THREADS 16
void *change_global_data(void *); /* for thr_create() */

class Mutected {
private:
    static mutex_t    Global_mutex;
    static int        Global_data;
public:
    static int        add_to_global_data(void);
    static int        subtract_from_global_data(void);
};

int Mutected::Global_data = 0;
mutex_t Mutected::Global_mutex;

int Mutected::add_to_global_data() {
    mutex_lock(&Global_mutex);
    Global_data++;
    mutex_unlock(&Global_mutex);
    return Global_data;
}

int Mutected::subtract_from_global_data() {
    mutex_lock(&Global_mutex);
    Global_data--;
    mutex_unlock(&Global_mutex);
    return Global_data;
}

void
main(int argc, char * argv[]) {
    int i=0;
    for (i=0;i< NUM_THREADS;i++) {
        thr_create(NULL, 0, change_global_data, NULL, 0, NULL);
    }
    while ((thr_join(NULL, NULL, NULL) == 0));
}

```

mutex_init(3C)

Interprocess Locking

```
void * change_global_data(void *)    {
    static int switcher = 0;
    if ((switcher++ % 3) == 0) /* one-in-three threads subtracts */
        cout << Mutected::subtract_from_global_data() << endl;
    else
        cout << Mutected::add_to_global_data() << endl;
    return NULL;
}
```

A mutex can protect data that is shared among processes. The mutex would need to be initialized as `USYNC_PROCESS`. One process initializes the process-shared mutex and writes it to a file to be mapped into memory by all cooperating processes (see `mmap(2)`). Afterwards, other independent processes can run the same program (whether concurrently or not) and share mutex-protected data.

```
/* cc thisfile.c -pthread */
/* To execute, run the command line "a.out 0 & a.out 1" */

#define _REENTRANT
#include <sys/types.h>
#include <sys/mman.h>
#include <sys/stat.h>
#include <fcntl.h>
#include <stdio.h>
#include <thread.h>
#define INTERPROCESS_FILE "ipc-sharedfile"
#define NUM_ADDTHREADS 12
#define NUM_SUBTRACTTHREADS 10
#define INCREMENT '0'
#define DECREMENT '1'
typedef struct {
    mutex_t      Interprocess_mutex;
    int          Interprocess_data;
} buffer_t;
buffer_t *buffer;

void *add_interprocess_data(), *subtract_interprocess_data();
void create_shared_memory(), test_argv();
int zeroed[sizeof(buffer_t)];
int ipc_fd, i=0;

void
main(int argc, char * argv[]){
    test_argv(argv[1]);

    switch (*argv[1]) {
    case INCREMENT:
        /* Initializes the process-shared mutex */
        /* Should be run prior to running a DECREMENT process */
        create_shared_memory();
        ipc_fd = open(INTERPROCESS_FILE, O_RDWR);
        buffer = (buffer_t *)mmap(NULL, sizeof(buffer_t),
            PROT_READ|PROT_WRITE, MAP_SHARED, ipc_fd, 0);
        buffer->Interprocess_data = 0;
        mutex_init(&buffer->Interprocess_mutex, USYNC_PROCESS, 0);
        for (i=0; i< NUM_ADDTHREADS; i++)
```

```

        thr_create(NULL, 0, add_interprocess_data, argv[1],
                  0, NULL);
        break;

    case DECREMENT:
        /* Should be run after the INCREMENT process has run. */
        while(ipc_fd = open(INTERPROCESS_FILE, O_RDWR) == -1)
            sleep(1);
        buffer = (buffer_t *)mmap(NULL, sizeof(buffer_t),
                                   PROT_READ|PROT_WRITE, MAP_SHARED, ipc_fd, 0);
        for (i=0; i< NUM_SUBTRACTTHREADS; i++)
            thr_create(NULL, 0, subtract_interprocess_data, argv[1],
                      0, NULL);
        break;
    } /* end switch */

    while ((thr_join(NULL, NULL, NULL) == 0));
} /* end main */

void *add_interprocess_data(char argv_1[]){
    mutex_lock(&buffer->Interprocess_mutex);
    buffer->Interprocess_data++;
    sleep(2);
    printf("%d is add-interprocess data, and %c is argv1\n",
           buffer->Interprocess_data, argv_1[0]);
    mutex_unlock(&buffer->Interprocess_mutex);
    return NULL;
}

void *subtract_interprocess_data(char argv_1[]) {
    mutex_lock(&buffer->Interprocess_mutex);
    buffer->Interprocess_data--;
    sleep(2);
    printf("%d is subtract-interprocess data, and %c is argv1\n",
           buffer->Interprocess_data, argv_1[0]);
    mutex_unlock(&buffer->Interprocess_mutex);
    return NULL;
}

void create_shared_memory(){
    int i;
    ipc_fd = creat(INTERPROCESS_FILE, O_CREAT|O_RDWR );
    for (i=0; i<sizeof(buffer_t); i++){
        zeroed[i] = 0;
        write(ipc_fd, &zeroed[i], 2);
    }
    close(ipc_fd);
    chmod(INTERPROCESS_FILE, S_IRWXU|S_IRWXG|S_IRWXO);
}

void test_argv(char argv1[]) {
    if (argv1 == NULL) {
        printf("use 0 as arg1 for initial process\n \
              or use 1 as arg1 for the second process\n");
        exit(NULL);
    }
}
}

```

mutex_init(3C)

Solaris Interprocess Robust Locking

A mutex can protect data that is shared among processes robustly. The mutex would need to be initialized as `USYNC_PROCESS_ROBUST`. One process initializes the robust process-shared mutex and writes it to a file to be mapped into memory by all cooperating processes (see `mmap(2)`). Afterwards, other independent processes can run the same program (whether concurrently or not) and share mutex-protected data.

The following example shows how to use a `USYNC_PROCESS_ROBUST` type mutex.

```
/* cc thisfile.c -pthread */
/* To execute, run the command line "a.out & a.out 1" */
#include <sys/types.h>
#include <sys/mman.h>
#include <fcntl.h>
#include <stdio.h>
#include <thread.h>
#define INTERPROCESS_FILE "ipc-sharedfile"
typedef struct {
    mutex_t    Interprocess_mutex;
    int        Interprocess_data;
} buffer_t;
buffer_t *buffer;
int make_data_consistent();
void create_shared_memory();
int zeroed[sizeof(buffer_t)];
int ipc_fd, i=0;
main(int argc, char * argv[]) {
    int rc;
    if (argc > 1) {
        while((ipc_fd = open(INTERPROCESS_FILE, O_RDWR)) == -1)
            sleep(1);
        buffer = (buffer_t *)mmap(NULL, sizeof(buffer_t),
            PROT_READ|PROT_WRITE, MAP_SHARED, ipc_fd, 0);
        mutex_init(&buffer->Interprocess_mutex,
            USYNC_PROCESS_ROBUST, 0);
    } else {
        create_shared_memory();
        ipc_fd = open(INTERPROCESS_FILE, O_RDWR);
        buffer = (buffer_t *)mmap(NULL, sizeof(buffer_t),
            PROT_READ|PROT_WRITE, MAP_SHARED, ipc_fd, 0);
        buffer->Interprocess_data = 0;
        mutex_init(&buffer->Interprocess_mutex,
            USYNC_PROCESS_ROBUST, 0);
    }
    for(;;) {
        rc = mutex_lock(&buffer->Interprocess_mutex);
        switch (rc) {
            case EOWNERDEAD:
                /* lock acquired.
                 * last owner died holding the lock, try to make
                 * the state associated with the mutex consistent.
                 * If so, make the robust lock consistent by
                 * reinitializing it.
                 */
                if (make_data_consistent())
                    mutex_init(&buffer->Interprocess_mutex,
                        USYNC_PROCESS_ROBUST, 0);
        }
    }
}
```

```

        mutex_unlock(&buffer->Interprocess_mutex);
        case ENOTRECOVERABLE:
            /* lock not acquired.
             * last owner got the mutex with EOWNERDEAD
             * mutex is not consistent (and data?),
             * so return from here
             */
            exit(1);
            break;
        case 0:
            /* no error - data is consistent */
            /* do something with data */
            mutex_unlock(&buffer->Interprocess_mutex);
            break;
    }
}
} /* end main */
void create_shared_memory() {
    int i;
    ipc_fd = creat(INTERPROCESS_FILE, O_CREAT|O_RDWR );
    for (i=0; i<sizeof(buffer_t); i++) {
        zeroed[i] = 0;
        write(ipc_fd, &zeroed[i], 2);
    }
    close(ipc_fd);
    chmod(INTERPROCESS_FILE, S_IRWXU|S_IRWXG|S_IRWXO);
}

/* return 1 if able to make data consistent, otherwise 0. */
int make_data_consistent () {
    buffer->Interprocess_data = 0;
    return (1);
}

```

Dynamically Allocated Mutexes

The following example allocates and frees memory in which a mutex is embedded.

```

struct record {
    int field1;
    int field2;
    mutex_t m;
} *r;
r = malloc(sizeof(struct record));
mutex_init(&r->m, USYNC_THREAD, NULL);
/*
 * The fields in this record are accessed concurrently
 * by acquiring the embedded lock.
 */

```

The thread execution in this example is as follows:

Thread 1 executes:

```

...
mutex_lock(&r->m);
r->field1++;
mutex_unlock(&r->m);
...

```

Thread 2 executes:

```

...
mutex_lock(&r->m);
localvar = r->field1;
mutex_unlock(&r->m);
...

```

mutex_init(3C)

Later, when a thread decides to free the memory pointed to by *r*, the thread should call `mutex_destroy()` on the mutexes in this memory.

In the following example, the main thread can do a `thr_join()` on both of the above threads. If there are no other threads using the memory in *r*, the main thread can now safely free *r*:

```
for (i = 0; i < 2; i++)
    thr_join(0, 0, 0);
mutex_destroy(&r->m); /* first destroy mutex */
free(r);             /* then free memory */
```

If the mutex is not destroyed, the program could have memory leaks.

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

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

SEE ALSO `mmap(2)`, `shmop(2)`, `pthread_mutex_init(3C)`, `pthread_mutexattr_settype(3C)`, `attributes(5)`, `mutex(5)`, `standards(5)`

NOTES The only supported policy is `SCHED_OTHER`. In Solaris under the `SCHED_OTHER` policy, there is no established order in which threads are unblocked.

The `mutex_lock()`, `mutex_unlock()`, and `mutex_trylock()` functions do not validate the mutex type. An uninitialized mutex or a mutex with an invalid type does not return `EINVAL`. Interfaces for mutexes with an invalid type have unspecified behavior.

Uninitialized mutexes that are allocated locally could contain junk data. Such mutexes need to be initialized using `mutex_init()`.

By default, if multiple threads are waiting for a mutex, the order of acquisition is undefined.

NAME	ndbm, dbm_clearerr, dbm_close, dbm_delete, dbm_error, dbm_fetch, dbm_firstkey, dbm_nextkey, dbm_open, dbm_store – database functions
SYNOPSIS	<pre>#include <ndbm.h> int dbm_clearerr (DBM *db) ; void dbm_close (DBM *db) ; int dbm_delete (DBM *db, datum key) ; int dbm_error (DBM *db) ; datum dbm_fetch (DBM *db, datum key) ; datum dbm_firstkey (DBM *db) ; datum dbm_nextkey (DBM *db) ; DBM *dbm_open (const char *file, int open_flags, mode_t file_mode) ; int dbm_store (DBM *db, datum key, datum content, int store_mode) ;</pre>
DESCRIPTION	<p>These functions create, access and modify a database. They maintain <i>key/content</i> pairs in a database. The functions will handle large databases (up to a billion blocks) and will access a keyed item in one or two file system accesses. This package replaces the earlier <code>dbm(3UCB)</code> library, which managed only a single database.</p> <p><i>keys</i> and <i>contents</i> are described by the datum typedef. A datum consists of at least two members, <i>dptr</i> and <i>dsize</i>. The <i>dptr</i> member points to an object that is <i>dsize</i> bytes in length. Arbitrary binary data, as well as ASCII character strings, may be stored in the object pointed to by <i>dptr</i>.</p> <p>The database is stored in two files. One file is a directory containing a bit map of keys and has <code>.dir</code> as its suffix. The second file contains all data and has <code>.pag</code> as its suffix.</p> <p>The <code>dbm_open()</code> function opens a database. The <code>file</code> argument to the function is the pathname of the database. The function opens two files named <code>file.dir</code> and <code>file.pag</code>. The <code>open_flags</code> argument has the same meaning as the <code>flags</code> argument of <code>open(2)</code> except that a database opened for write-only access opens the files for read and write access. The <code>file_mode</code> argument has the same meaning as the third argument of <code>open(2)</code>.</p> <p>The <code>dbm_close()</code> function closes a database. The argument <code>db</code> must be a pointer to a <code>dbm</code> structure that has been returned from a call to <code>dbm_open()</code>.</p> <p>The <code>dbm_fetch()</code> function reads a record from a database. The argument <code>db</code> is a pointer to a database structure that has been returned from a call to <code>dbm_open()</code>. The argument <code>key</code> is a datum that has been initialized by the application program to the value of the key that matches the key of the record the program is fetching.</p> <p>The <code>dbm_store()</code> function writes a record to a database. The argument <code>db</code> is a pointer to a database structure that has been returned from a call to <code>dbm_open()</code>. The argument <code>key</code> is a datum that has been initialized by the application program to the</p>

value of the key that identifies (for subsequent reading, writing or deleting) the record the program is writing. The argument *content* is a datum that has been initialized by the application program to the value of the record the program is writing. The argument *store_mode* controls whether `dbm_store()` replaces any pre-existing record that has the same key that is specified by the *key* argument. The application program must set *store_mode* to either `DBM_INSERT` or `DBM_REPLACE`. If the database contains a record that matches the *key* argument and *store_mode* is `DBM_REPLACE`, the existing record is replaced with the new record. If the database contains a record that matches the *key* argument and *store_mode* is `DBM_INSERT`, the existing record is not replaced with the new record. If the database does not contain a record that matches the *key* argument and *store_mode* is either `DBM_INSERT` or `DBM_REPLACE`, the new record is inserted in the database.

The `dbm_delete()` function deletes a record and its key from the database. The argument *db* is a pointer to a database structure that has been returned from a call to `dbm_open()`. The argument *key* is a datum that has been initialized by the application program to the value of the key that identifies the record the program is deleting.

The `dbm_firstkey()` function returns the first key in the database. The argument *db* is a pointer to a database structure that has been returned from a call to `dbm_open()`.

The `dbm_nextkey()` function returns the next key in the database. The argument *db* is a pointer to a database structure that has been returned from a call to `dbm_open()`. The `dbm_firstkey()` function must be called before calling `dbm_nextkey()`. Subsequent calls to `dbm_nextkey()` return the next key until all of the keys in the database have been returned.

The `dbm_error()` function returns the error condition of the database. The argument *db* is a pointer to a database structure that has been returned from a call to `dbm_open()`.

The `dbm_clearerr()` function clears the error condition of the database. The argument *db* is a pointer to a database structure that has been returned from a call to `dbm_open()`.

These database functions support key/content pairs of at least 1024 bytes.

RETURN VALUES

The `dbm_store()` and `dbm_delete()` functions return 0 when they succeed and a negative value when they fail.

The `dbm_store()` function returns 1 if it is called with a *flags* value of `DBM_INSERT` and the function finds an existing record with the same key.

The `dbm_error()` function returns 0 if the error condition is not set and returns a non-zero value if the error condition is set.

The return value of `dbm_clearerr()` is unspecified.

The `dbm_firstkey()` and `dbm_nextkey()` functions return a key datum. When the end of the database is reached, the `dptr` member of the key is a null pointer. If an error is detected, the `dptr` member of the key is a null pointer and the error condition of the database is set.

The `dbm_fetch()` function returns a content datum. If no record in the database matches the key or if an error condition has been detected in the database, the `dptr` member of the content is a null pointer.

The `dbm_open()` function returns a pointer to a database structure. If an error is detected during the operation, `dbm_open()` returns a `(DBM *)0`.

ERRORS No errors are defined.

USAGE The following code can be used to traverse the database:

```
for(key = dbm_firstkey(db); key.dptr != NULL; key = dbm_nextkey(db))
```

The `dbm_*` functions provided in this library should not be confused in any way with those of a general-purpose database management system. These functions do not provide for multiple search keys per entry, they do not protect against multi-user access (in other words they do not lock records or files), and they do not provide the many other useful database functions that are found in more robust database management systems. Creating and updating databases by use of these functions is relatively slow because of data copies that occur upon hash collisions. These functions are useful for applications requiring fast lookup of relatively static information that is to be indexed by a single key.

The `dptr` pointers returned by these functions may point into static storage that may be changed by subsequent calls.

The `dbm_delete()` function does not physically reclaim file space, although it does make it available for reuse.

After calling `dbm_store()` or `dbm_delete()` during a pass through the keys by `dbm_firstkey()` and `dbm_nextkey()`, the application should reset the database by calling `dbm_firstkey()` before again calling `dbm_nextkey()`.

EXAMPLES **EXAMPLE 1** Using the Database Functions

The following example stores and retrieves a phone number, using the name as the key. Note that this example does not include error checking.

```
#include <ndbm.h>
#include <stdio.h>
#include <fcntl.h>
#define NAME "Bill"
#define PHONE_NO "123-4567"
#define DB_NAME "phones"
main()
{
    DBM *db;
    datum name = {NAME, sizeof (NAME)};
```

EXAMPLE 1 Using the Database Functions (Continued)

```

datum put_phone_no = {PHONE_NO, sizeof (PHONE_NO)};
datum get_phone_no;
/* Open the database and store the record */
db = dbm_open(DB_NAME, O_RDWR | O_CREAT, 0660);
(void) dbm_store(db, name, put_phone_no, DBM_INSERT);
/* Retrieve the record */
get_phone_no = dbm_fetch(db, name);
(void) printf("Name: %s, Phone Number: %s\n", name.dptr,
get_phone_no.dptr);
/* Close the database */
dbm_close(db);
return (0);
}

```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Unsafe

SEE ALSO `ar(1)`, `cat(1)`, `cp(1)`, `tar(1)`, `open(2)`, `dbm(3UCB)`, `netconfig(4)`, `attributes(5)`, `standards(5)`

NOTES The `.pag` file will contain holes so that its apparent size may be larger than its actual content. Older versions of the UNIX operating system may create real file blocks for these holes when touched. These files cannot be copied by normal means (`cp(1)`, `cat(1)`, `tar(1)`, `ar(1)`) without filling in the holes.

The sum of the sizes of a *key/content* pair must not exceed the internal block size (currently 1024 bytes). Moreover all *key/content* pairs that hash together must fit on a single block. `dbm_store()` will return an error in the event that a disk block fills with inseparable data.

The order of keys presented by `dbm_firstkey()` and `dbm_nextkey()` depends on a hashing function.

There are no interlocks and no reliable cache flushing; thus concurrent updating and reading is risky.

The database files (`file.dir` and `file.pag`) are binary and are architecture-specific (for example, they depend on the architecture's byte order.) These files are not guaranteed to be portable across architectures.

NAME	nice – change priority of a process
SYNOPSIS	<pre> /usr/ucb/cc [flag ...] file ... #include<unistd.h> int nice (incr) ; int incr; </pre>
DESCRIPTION	<p>The scheduling priority of the process is augmented by <i>incr</i>. Positive priorities get less service than normal. Priority 10 is recommended to users who wish to execute long-running programs without undue impact on system performance.</p> <p>Negative increments are illegal, except when specified by the privileged user. The priority is limited to the range -20 (most urgent) to 20 (least). Requests for values above or below these limits result in the scheduling priority being set to the corresponding limit.</p> <p>The priority of a process is passed to a child process by <code>fork(2)</code>. For a privileged process to return to normal priority from an unknown state, <code>nice()</code> should be called successively with arguments -40 (goes to priority -20 because of truncation), 20 (to get to 0), then 0 (to maintain compatibility with previous versions of this call).</p>
RETURN VALUES	Upon successful completion, <code>nice()</code> returns 0. Otherwise, a value of -1 is returned and <code>errno</code> is set to indicate the error.
ERRORS	<p>The priority is not changed if:</p> <p><code>EPERM</code> The value of <i>incr</i> specified was negative, and the effective user ID is not the privileged user.</p>
SEE ALSO	<code>nice(1)</code> , <code>renice(1)</code> , <code>fork(2)</code> , <code>priocntl(2)</code> , <code>getpriority(3C)</code>
NOTES	Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-threaded applications is unsupported.

nlist(3UCB)

NAME	nlist – get entries from symbol table
SYNOPSIS	<pre>/usr/ucb/cc [flag ...] file ... #include <nlist.h> int nlist (filename, nl) ; char *filename ; struct nlist *nl ;</pre>
DESCRIPTION	<p>nlist() examines the symbol table from the executable image whose name is pointed to by <i>filename</i>, and selectively extracts a list of values and puts them in the array of nlist structures pointed to by <i>nl</i>. The name list pointed to by <i>nl</i> consists of an array of structures containing names, types and values. The n_name field of each such structure is taken to be a pointer to a character string representing a symbol name. The list is terminated by an entry with a NULL pointer (or a pointer to a NULL string) in the n_name field. For each entry in <i>nl</i>, if the named symbol is present in the executable image's symbol table, its value and type are placed in the n_value and n_type fields. If a symbol cannot be located, the corresponding n_type field of <i>nl</i> is set to zero.</p>
RETURN VALUES	Upon normal completion, nlist() returns the number of symbols that were not located in the symbol table. If an error occurs, nlist() returns -1 and sets all of the n_type fields in members of the array pointed to by <i>nl</i> to zero.
SEE ALSO	nlist(3ELF), a.out(4)
NOTES	<p>Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.</p> <p>Only the n_value field is compatibly set. Other fields in the nlist structure are filled with the ELF (Executable and Linking Format) values (see nlist(3ELF) and a.out(4)).</p>

NAME	nl_langinfo – language information								
SYNOPSIS	<pre>#include <langinfo.h> char *nl_langinfo(nl_item item);</pre>								
DESCRIPTION	<p>The <code>nl_langinfo()</code> function returns a pointer to a null-terminated string containing information relevant to a particular language or cultural area defined in the programs locale. The manifest constant names and values of <i>item</i> are defined by <code><langinfo.h></code>. For example:</p> <pre>nl_langinfo (ABDAY_1);</pre> <p>would return a pointer to the string “Dim” if the identified language was French and a French locale was correctly installed; or “Sun” if the identified language was English.</p>								
RETURN VALUES	<p>If <code>setlocale(3C)</code> has not been called successfully, or if data for a supported language is either not available, or if <i>item</i> is not defined therein, then <code>nl_langinfo()</code> returns a pointer to the corresponding string in the C locale. In all locales, <code>nl_langinfo()</code> returns a pointer to an empty string if <i>item</i> contains an invalid setting.</p>								
USAGE	<p>The <code>nl_langinfo()</code> function can be used safely in multithreaded applications, as long as <code>setlocale(3C)</code> is not being called to change the locale.</p>								
ATTRIBUTES	<p>See <code>attributes(5)</code> 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>CSI</td> <td>Enabled</td> </tr> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	Standard	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
CSI	Enabled								
Interface Stability	Standard								
MT-Level	MT-Safe with exceptions								
SEE ALSO	<p><code>setlocale(3C)</code>, <code>langinfo.h(3HEAD)</code>, <code>nl_types.h(3HEAD)</code>, <code>attributes(5)</code>, <code>standards(5)</code></p>								
WARNINGS	<p>The array pointed to by the return value should not be modified by the program. Subsequent calls to <code>nl_langinfo()</code> may overwrite the array.</p>								

offsetof(3C)

NAME	offsetof – offset of structure member				
SYNOPSIS	<pre>#include <stddef.h> size_t offsetof(<i>type</i>, <i>member-designator</i>);</pre>				
DESCRIPTION	The <code>offsetof()</code> macro defined in <code><stddef.h></code> expands to an integral constant expression that has type <code>size_t</code> . The value of this expression is the offset in bytes to the structure member (designated by <i>member-designator</i>) from the beginning of its structure (designated by <i>type</i>).				
ATTRIBUTES	See <code>attributes(5)</code> 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>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>attributes(5)</code>				

NAME	opendir, fdopendir – open directory																
SYNOPSIS	<pre>#include <sys/types.h> #include <dirent.h> DIR *opendir (const char *dirname) ; DIR *fdopendir (int fildes) ;</pre>																
DESCRIPTION	<p>The <code>opendir()</code> function opens a directory stream corresponding to the directory named by the <i>dirname</i> argument.</p> <p>The <code>fdopendir()</code> function opens a directory stream for the directory file descriptor <i>fildes</i>. The directory file descriptor should not be used or closed following a successful function call, as this might cause undefined results from future operations on the directory stream obtained from the call. Use <code>closedir(3C)</code> to close a directory stream.</p> <p>The directory stream is positioned at the first entry. If the type <code>DIR</code> is implemented using a file descriptor, applications will only be able to open up to a total of <code>{OPEN_MAX}</code> files and directories. A successful call to any of the <code>exec</code> functions will close any directory streams that are open in the calling process. See <code>exec(2)</code>.</p>																
RETURN VALUES	Upon successful completion, <code>opendir()</code> and <code>fdopendir()</code> return a pointer to an object of type <code>DIR</code> . Otherwise, a null pointer is returned and <code>errno</code> is set to indicate the error.																
ERRORS	<p>The <code>opendir()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>EACCES</code></td> <td>Search permission is denied for the component of the path prefix of <i>dirname</i> or read permission is denied for <i>dirname</i>.</td> </tr> <tr> <td><code>ELOOP</code></td> <td>Too many symbolic links were encountered in resolving <i>path</i>.</td> </tr> <tr> <td><code>ENAMETOOLONG</code></td> <td>The length of the <i>dirname</i> argument exceeds <code>{PATH_MAX}</code>, or a path name component is longer than <code>{NAME_MAX}</code> while <code>{_POSIX_NO_TRUNC}</code> is in effect.</td> </tr> <tr> <td><code>ENOENT</code></td> <td>A component of <i>dirname</i> does not name an existing directory or <i>dirname</i> is an empty string.</td> </tr> <tr> <td><code>ENOTDIR</code></td> <td>A component of <i>dirname</i> is not a directory.</td> </tr> </table> <p>The <code>fdopendir()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>ENOTDIR</code></td> <td>The file descriptor <i>fildes</i> does not reference a directory.</td> </tr> </table> <p>The <code>opendir()</code> function may fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>EMFILE</code></td> <td>There are <code>{OPEN_MAX}</code> file descriptors currently open in the calling process.</td> </tr> <tr> <td><code>ENAMETOOLONG</code></td> <td>Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>PATH_MAX</code>.</td> </tr> </table>	<code>EACCES</code>	Search permission is denied for the component of the path prefix of <i>dirname</i> or read permission is denied for <i>dirname</i> .	<code>ELOOP</code>	Too many symbolic links were encountered in resolving <i>path</i> .	<code>ENAMETOOLONG</code>	The length of the <i>dirname</i> argument exceeds <code>{PATH_MAX}</code> , or a path name component is longer than <code>{NAME_MAX}</code> while <code>{_POSIX_NO_TRUNC}</code> is in effect.	<code>ENOENT</code>	A component of <i>dirname</i> does not name an existing directory or <i>dirname</i> is an empty string.	<code>ENOTDIR</code>	A component of <i>dirname</i> is not a directory.	<code>ENOTDIR</code>	The file descriptor <i>fildes</i> does not reference a directory.	<code>EMFILE</code>	There are <code>{OPEN_MAX}</code> file descriptors currently open in the calling process.	<code>ENAMETOOLONG</code>	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>PATH_MAX</code> .
<code>EACCES</code>	Search permission is denied for the component of the path prefix of <i>dirname</i> or read permission is denied for <i>dirname</i> .																
<code>ELOOP</code>	Too many symbolic links were encountered in resolving <i>path</i> .																
<code>ENAMETOOLONG</code>	The length of the <i>dirname</i> argument exceeds <code>{PATH_MAX}</code> , or a path name component is longer than <code>{NAME_MAX}</code> while <code>{_POSIX_NO_TRUNC}</code> is in effect.																
<code>ENOENT</code>	A component of <i>dirname</i> does not name an existing directory or <i>dirname</i> is an empty string.																
<code>ENOTDIR</code>	A component of <i>dirname</i> is not a directory.																
<code>ENOTDIR</code>	The file descriptor <i>fildes</i> does not reference a directory.																
<code>EMFILE</code>	There are <code>{OPEN_MAX}</code> file descriptors currently open in the calling process.																
<code>ENAMETOOLONG</code>	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>PATH_MAX</code> .																

opendir(3C)

ENFILE Too many files are currently open on the system.

USAGE The `opendir()` and `fdopendir()` functions should be used in conjunction with `readdir(3C)`, `closedir(3C)` and `rewinddir(3C)` to examine the contents of the directory (see the `EXAMPLES` section in `readdir(3C)`). This method is recommended for portability.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	<code>opendir()</code> is Standard; <code>fdopendir()</code> is Evolving
MT-Level	Safe

SEE ALSO `lstat(2)`, `symlink(2)`, `closedir(3C)`, `readdir(3C)`, `rewinddir(3C)`, `attributes(5)`

NAME	perror, errno – print system error messages						
SYNOPSIS	<pre>#include <stdio.h> void perror(const char *s); #include <errno.h> int errno;</pre>						
DESCRIPTION	<p>The <code>perror()</code> function produces a message on the standard error output (file descriptor 2) describing the last error encountered during a call to a system or library function. The argument string <code>s</code> is printed, followed by a colon and a blank, followed by the message and a NEWLINE character. If <code>s</code> is a null pointer or points to a null string, the colon is not printed. The argument string should include the name of the program that incurred the error. The error number is taken from the external variable <code>errno</code>, which is set when errors occur but not cleared when non-erroneous calls are made. See intro(2).</p> <p>In the case of multithreaded applications, the <code>-mt</code> option must be specified on the command line at compilation time (see threads(5)). When the <code>-mt</code> option is specified, <code>errno</code> becomes a macro that enables each thread to have its own <code>errno</code>. This <code>errno</code> macro can be used on either side of the assignment as though it were a variable.</p>						
USAGE	If the application is linked with <code>-lintl</code> , then messages printed from this function are in the native language specified by the <code>LC_MESSAGES</code> locale category. See setlocale(3C) .						
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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	intro(2) , fmtmsg(3C) , gettext(3C) , setlocale(3C) , strerror(3C) , attributes(5) , standards(5) , threads(5)						

pfmt(3C)

NAME	pfmt – display error message in standard format
SYNOPSIS	<pre>#include <pfmt.h> int pfmt(FILE *stream, long flags, char *format, ... /* arg */);</pre>
DESCRIPTION	<p>The <code>pfmt()</code> retrieves a format string from a locale-specific message database (unless <code>MM_NOGET</code> is specified) and uses it for <code>printf(3C)</code> style formatting of <code>args</code>. The output is displayed on <code>stream</code>.</p> <p>The <code>pfmt()</code> function encapsulates the output in the standard error message format (unless <code>MM_NOSTD</code> is specified, in which case the output is similar to <code>printf()</code>).</p> <p>If the <code>printf()</code> format string is to be retrieved from a message database, the <code>format</code> argument must have the following structure:</p> <pre><catalog> : <msgnum> : <defmsg>.</pre> <p>If <code>MM_NOGET</code> is specified, only the <code>defmsg</code> field must be specified.</p> <p>The <code>catalog</code> field is used to indicate the message database that contains the localized version of the format string. This field must be limited to 14 characters selected from the set of all characters values, excluding <code>\0</code> (null) and the ASCII codes for <code>/</code> (slash) and <code>:</code> (colon).</p> <p>The <code>msgnum</code> field is a positive number that indicates the index of the string into the message database.</p> <p>If the catalog does not exist in the locale (specified by the last call to <code>setlocale(3C)</code> using the <code>LC_ALL</code> or <code>LC_MESSAGES</code> categories), or if the message number is out of bound, <code>pfmt()</code> will attempt to retrieve the message from the C locale. If this second retrieval fails, <code>pfmt()</code> uses the <code>defmsg</code> field of the <code>format</code> argument.</p> <p>If <code>catalog</code> is omitted, <code>pfmt()</code> will attempt to retrieve the string from the default catalog specified by the last call to <code>setcat(3C)</code>. In this case, the <code>format</code> argument has the following structure:</p> <pre>: <msgnum> : <defmsg>.</pre> <p>The <code>pfmt()</code> will output <code>Message not found!\n</code> as format string if <code>catalog</code> is not a valid catalog name, if no catalog is specified (either explicitly or with <code>setcat()</code>), if <code>msgnum</code> is not a valid number, or if no message could be retrieved from the message databases and <code>defmsg</code> was omitted.</p> <p>The <code>flags</code> argument determine the type of output (such as whether the <code>format</code> should be interpreted as is or encapsulated in the standard message format), and the access to message catalogs to retrieve a localized version of <code>format</code>.</p> <p>The <code>flags</code> argument is composed of several groups, and can take the following values (one from each group):</p>

Output format control

MM_NOSTD	Do not use the standard message format, interpret <code>format</code> as <code>printf()</code> format. Only <i>catalog access control flags</i> should be specified if MM_NOSTD is used; all other flags will be ignored.
MM_STD	Output using the standard message format (default value 0).

Catalog access control

MM_NOGET	Do not retrieve a localized version of <code>format</code> . In this case, only the <i>defmsg</i> field of the <code>format</code> is specified.
MM_GET	Retrieve a localized version of <code>format</code> from the <i>catalog</i> , using <i>msgid</i> as the index and <i>defmsg</i> as the default message (default value 0).

Severity (standard message format only)

MM_HALT	Generate a localized version of HALT, but do not halt the machine.
MM_ERROR	Generate a localized version of ERROR (default value 0).
MM_WARNING	Generate a localized version of WARNING.
MM_INFO	Generate a localized version of INFO.

Additional severities can be defined. Add-on severities can be defined with number-string pairs with numeric values from the range [5-255], using [addsev\(3C\)](#). The specified severity will be generated from the bitwise OR operation of the numeric value and other *flags*. If the severity is not defined, `pfmt()` uses the string `SEV=N`, where *N* is replaced by the integer severity value passed in *flags*.

Multiple severities passed in *flags* will not be detected as an error. Any combination of severities will be summed and the numeric value will cause the display of either a severity string (if defined) or the string `SEV=N` (if undefined).

Action

MM_ACTION	Specify an action message. Any severity value is superseded and replaced by a localized version of TO FIX.
-----------	--

**STANDARD
ERROR
MESSAGE
FORMAT**

The `pfmt()` function displays error messages in the following format:

label: *severity*: *text*

If no *label* was defined by a call to [setLabel\(3C\)](#), the message is displayed in the format:

severity: *text*

If `pfmt()` is called twice to display an error message and a helpful *action* or recovery message, the output can look like:

pfmt(3C)

label: severity: textlabel: TO FIX: text

RETURN VALUES

Upon success, `pfmt()` returns the number of bytes transmitted. Upon failure, it returns a negative value:

-1 Write error to *stream*.

EXAMPLES

EXAMPLE 1 Example of `pfmt()` function.

Example 1:

```
setlabel("UX:test");
pfmt(stderr, MM_ERROR, "test:2:Cannot open file: %s\n", strerror(errno));
```

displays the message:

```
UX:test: ERROR: Cannot open file: No such file or directory
```

Example 2:

```
setlabel("UX:test");
setcat("test");
pfmt(stderr, MM_ERROR, ":10:Syntax error\n");
pfmt(stderr, MM_ACTION, "55:Usage ... \n");
```

displays the message

```
UX:test: ERROR: Syntax error
UX:test: TO FIX: Usage ...
```

USAGE

Since it uses [gettext\(3C\)](#), `pfmt()` should not be used.

ATTRIBUTES

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-safe

SEE ALSO

[addsev\(3C\)](#), [gettext\(3C\)](#), [lfmt\(3C\)](#), [printf\(3C\)](#), [setcat\(3C\)](#), [setlabel\(3C\)](#), [setlocale\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#)

NAME	plock – lock or unlock into memory process, text, or data								
SYNOPSIS	<pre>#include <sys/lock.h> int plock(int op);</pre>								
DESCRIPTION	<p>The plock() function allows the calling process to lock or unlock into memory its text segment (text lock), its data segment (data lock), or both its text and data segments (process lock). Locked segments are immune to all routine swapping. The effective user ID of the calling process must be super-user to use this call.</p> <p>The plock() function performs the function specified by <i>op</i>:</p> <table border="0"> <tr> <td>PROCLCK</td> <td>Lock text and data segments into memory (process lock).</td> </tr> <tr> <td>TXTLCK</td> <td>Lock text segment into memory (text lock).</td> </tr> <tr> <td>DATLCK</td> <td>Lock data segment into memory (data lock).</td> </tr> <tr> <td>UNLCK</td> <td>Remove locks.</td> </tr> </table>	PROCLCK	Lock text and data segments into memory (process lock).	TXTLCK	Lock text segment into memory (text lock).	DATLCK	Lock data segment into memory (data lock).	UNLCK	Remove locks.
PROCLCK	Lock text and data segments into memory (process lock).								
TXTLCK	Lock text segment into memory (text lock).								
DATLCK	Lock data segment into memory (data lock).								
UNLCK	Remove locks.								
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 plock() function fails and does not perform the requested operation if:</p> <table border="0"> <tr> <td>EAGAIN</td> <td>Not enough memory.</td> </tr> <tr> <td>EINVAL</td> <td>The <i>op</i> argument is equal to PROCLCK and a process lock, a text lock, or a data lock already exists on the calling process; the <i>op</i> argument is equal to TXTLCK and a text lock or a process lock already exists on the calling process; the <i>op</i> argument is equal to DATLCK and a data lock or a process lock already exists on the calling process; or the <i>op</i> argument is equal to UNLCK and no lock exists on the calling process.</td> </tr> <tr> <td>EPERM</td> <td>The {PRIV_PROC_LOCK_MEMORY} privilege is not asserted in the effective set of the calling process.</td> </tr> </table>	EAGAIN	Not enough memory.	EINVAL	The <i>op</i> argument is equal to PROCLCK and a process lock, a text lock, or a data lock already exists on the calling process; the <i>op</i> argument is equal to TXTLCK and a text lock or a process lock already exists on the calling process; the <i>op</i> argument is equal to DATLCK and a data lock or a process lock already exists on the calling process; or the <i>op</i> argument is equal to UNLCK and no lock exists on the calling process.	EPERM	The {PRIV_PROC_LOCK_MEMORY} privilege is not asserted in the effective set of the calling process.		
EAGAIN	Not enough memory.								
EINVAL	The <i>op</i> argument is equal to PROCLCK and a process lock, a text lock, or a data lock already exists on the calling process; the <i>op</i> argument is equal to TXTLCK and a text lock or a process lock already exists on the calling process; the <i>op</i> argument is equal to DATLCK and a data lock or a process lock already exists on the calling process; or the <i>op</i> argument is equal to UNLCK and no lock exists on the calling process.								
EPERM	The {PRIV_PROC_LOCK_MEMORY} privilege is not asserted in the effective set of the calling process.								
USAGE	The <code>mlock(3C)</code> and <code>mlockall(3C)</code> functions are the preferred interfaces for process locking.								
ATTRIBUTES	See <code>attributes(5)</code> 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>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe				
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
MT-Level	MT-Safe								
SEE ALSO	<code>exec(2)</code> , <code>exit(2)</code> , <code>fork(2)</code> , <code>memcntl(2)</code> , <code>mlock(3C)</code> , <code>mlockall(3C)</code> , <code>attributes(5)</code>								

popen(3C)

NAME	popen, pclose – initiate a pipe to or from a process						
SYNOPSIS	<pre>#include <stdio.h> FILE *popen(const char *command, const char *mode); int pclose(FILE *stream);</pre>						
DESCRIPTION	<p>The popen() function creates a pipe between the calling program and the command to be executed. The arguments to popen() are pointers to null-terminated strings. The <i>command</i> argument consists of a shell command line. The <i>mode</i> argument is an I/O mode, either <i>r</i> for reading or <i>w</i> for writing. The value returned is a stream pointer such that one can write to the standard input of the command, if the I/O mode is <i>w</i>, by writing to the file <i>stream</i> (see intro(3)); and one can read from the standard output of the command, if the I/O mode is <i>r</i>, by reading from the file <i>stream</i>. Because open files are shared, a type <i>r</i> command may be used as an input filter and a type <i>w</i> as an output filter.</p> <p>The environment of the executed command will be as if a child process were created within the popen() call using fork(2). If the application is standard-conforming (see standards(5)), the child is invoked with the call:</p> <pre>execl("/usr/xpg4/bin/sh", "sh", "-c", command, (char *)0);</pre> otherwise, the child is invoked with the call: <pre>execl("/usr/bin/sh", "sh", "-c", command, (char *)0);</pre> <p>The pclose() function closes a stream opened by popen() by closing the pipe. It waits for the associated process to terminate and returns the termination status of the process running the command language interpreter. This is the value returned by waitpid(3C). See wait.h(3HEAD) for more information on termination status.</p>						
RETURN VALUES	<p>Upon successful completion, popen() returns a pointer to an open stream that can be used to read or write to the pipe. Otherwise, it returns a null pointer and may set errno to indicate the error.</p> <p>Upon successful completion, pclose() returns the termination status of the command language interpreter as returned by waitpid(). Otherwise, it returns -1 and sets errno to indicate the error.</p>						
ERRORS	<p>The pclose() function will fail if:</p> <table><tr><td>ECHILD</td><td>The status of the child process could not be obtained, as described in the DESCRIPTION.</td></tr></table> <p>The popen() function may fail if:</p> <table><tr><td>EMFILE</td><td>There are currently FOPEN_MAX or STREAM_MAX streams open in the calling process.</td></tr><tr><td>EINVAL</td><td>The mode argument is invalid.</td></tr></table>	ECHILD	The status of the child process could not be obtained, as described in the DESCRIPTION.	EMFILE	There are currently FOPEN_MAX or STREAM_MAX streams open in the calling process.	EINVAL	The mode argument is invalid.
ECHILD	The status of the child process could not be obtained, as described in the DESCRIPTION.						
EMFILE	There are currently FOPEN_MAX or STREAM_MAX streams open in the calling process.						
EINVAL	The mode argument is invalid.						

The `popen()` function may also set `errno` values as described by `fork(2)` or `pipe(2)`.

USAGE If the original and `popen()` processes concurrently read or write a common file, neither should use buffered I/O. Problems with an output filter may be forestalled by careful buffer flushing, for example, with `fflush()` (see `fclose(3C)`). A security hole exists through the `IFS` and `PATH` environment variables. Full pathnames should be used (or `PATH` reset) and `IFS` should be set to space and tab ("`\t`").

The signal handler for `SIGCHLD` should be set to default when using `popen()`. If the process has established a signal handler for `SIGCHLD`, it will be called when the command terminates. If the signal handler or another thread in the same process issues a `wait(3C)` call, it will interfere with the return value of `pclose()`. If the process's signal handler for `SIGCHLD` has been set to ignore the signal, `pclose()` will fail and `errno` will be set to `ECHILD`.

EXAMPLES **EXAMPLE 1** `popen()` example

The following program will print on the standard output (see `stdio(3C)`) the names of files in the current directory with a `.c` suffix.

```
#include <stdio.h>
#include <stdlib.h>
main( )
{
    char *cmd = "/usr/bin/ls *.c";
    char buf[BUFSIZ];
    FILE *ptr;

    if ((ptr = popen(cmd, "r")) != NULL)
        while (fgets(buf, BUFSIZ, ptr) != NULL)
            (void) printf("%s", buf);
        (void) pclose(ptr);
    return 0;
}
```

EXAMPLE 2 `system()` replacement

The following function can be used in a multithreaded process in place of the most common usage of the Unsafe `system(3C)` function:

```
int my_system(const char *cmd)
{
    FILE *p;

    if ((p = popen(cmd, "w")) == NULL)
        return (-1);
    return (pclose(p));
}
```

popen(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Safe

SEE ALSO `ksh(1)`, `pipe(2)`, `fclose(3C)`, `fopen(3C)`, `stdio(3C)`, `system(3C)`, `wait(3C)`, `waitpid(3C)`, `wait.h(3HEAD)`, `attributes(5)`, `standards(5)`

NAME	port_alert – set a port in alert mode								
SYNOPSIS	<pre>#include <port.h> int port_alert(int port, int flags, int events, void *user);</pre>								
DESCRIPTION	<p>The <code>port_alert()</code> function transitions a port into or out of alert mode. A port in alert mode immediately awakens all threads blocked in <code>port_get(3C)</code> or <code>port_getn(3C)</code>. These threads return with an alert notification that consists of a single <code>port_event_t</code> structure with the source <code>PORT_SOURCE_ALERT</code>. Subsequent threads trying to retrieve events from a port that is in alert mode will return immediately with the alert notification.</p> <p>A port is transitioned into alert mode by calling the <code>port_alert()</code> function with a non-zero <code>events</code> parameter. The specified <code>events</code> and <code>user</code> parameters will be made available in the <code>portev_events</code> and the <code>portev_user</code> members of the alert notification, respectively. The <code>flags</code> argument determines the mode of operation of the alert mode:</p> <ul style="list-style-type: none"> ■ If <code>flags</code> is set to <code>PORT_ALERT_SET</code>, <code>port_alert()</code> sets the port in alert mode independent of the current state of the port. The <code>portev_events</code> and <code>portev_user</code> members are set or updated accordingly. ■ If <code>flags</code> is set to <code>PORT_ALERT_UPDATE</code> and the port is not in alert mode, <code>port_alert()</code> transitions the port into alert mode. The <code>portev_events</code> and <code>portev_user</code> members are set accordingly. ■ If <code>flags</code> is set to <code>PORT_ALERT_UPDATE</code> and the port is already in alert mode, <code>port_alert()</code> returns with an error value of <code>EBUSY</code>. <p><code>PORT_ALERT_SET</code> and <code>PORT_ALERT_UPDATE</code> are mutually exclusive.</p> <p>A port is transitioned out of alert mode by calling the <code>port_alert()</code> function with a zero <code>events</code> parameter.</p> <p>Events can be queued to a port that is in alert mode, but they will not be retrievable until the port is transitioned out of alert mode.</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 <code>port_alert()</code> function will fail if:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;"><code>EBADF</code></td> <td>The port identifier is not valid.</td> </tr> <tr> <td><code>EBADFD</code></td> <td>The <code>port</code> argument is not an event port file descriptor.</td> </tr> <tr> <td><code>EBUSY</code></td> <td>The port is already in alert mode.</td> </tr> <tr> <td><code>EINVAL</code></td> <td>Mutually exclusive <code>flags</code> are set.</td> </tr> </table>	<code>EBADF</code>	The port identifier is not valid.	<code>EBADFD</code>	The <code>port</code> argument is not an event port file descriptor.	<code>EBUSY</code>	The port is already in alert mode.	<code>EINVAL</code>	Mutually exclusive <code>flags</code> are set.
<code>EBADF</code>	The port identifier is not valid.								
<code>EBADFD</code>	The <code>port</code> argument is not an event port file descriptor.								
<code>EBUSY</code>	The port is already in alert mode.								
<code>EINVAL</code>	Mutually exclusive <code>flags</code> are set.								

port_alert(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	all
Availability	SUNWcsr, SUNWhea
Interface Stability	Evolving
MT-Level	Safe

SEE ALSO `port_associate(3C)`, `port_create(3C)`, `port_get(3C)`, `port_send(3C)`, `attributes(5)`

NAME	port_associate, port_dissociate – associate or dissociate the object with the port						
SYNOPSIS	<pre>#include <port.h> int port_associate(int port, int source, uintptr_t object, int events, void *user); int port_dissociate(int port, int source, uintptr_t object);</pre>						
DESCRIPTION	<p>The <code>port_associate()</code> function associates specific <i>events</i> of a given <i>object</i> with a <i>port</i>. Only objects associated with a particular port are able to generate events that can be retrieved using <code>port_get(3C)</code> or <code>port_getn(3C)</code>. The delivery event has its <code>portev_user</code> member set to the value specified in the <i>user</i> parameter. If the specified object is already associated with the specified port, the <code>port_associate()</code> function serves to update the <i>events</i> and <i>user</i> arguments of the association. The <code>port_dissociate()</code> function removes the association of an object with a port.</p> <p>The only objects associated with a port by way of the <code>port_associate()</code> function are objects of type <code>PORT_SOURCE_FD</code>. Objects of other types have type-specific association mechanisms. See <code>port_create(3C)</code> for details.</p> <p>Objects of type <code>PORT_SOURCE_FD</code> are file descriptors. The event types for <code>PORT_SOURCE_FD</code> objects are described in <code>poll(2)</code>. At most one event notification will be generated per associated file descriptor. For example, if a file descriptor is associated with a port for the <code>POLLRDNORM</code> event and data is available on the file descriptor at the time the <code>port_associate()</code> function is called, an event is immediately sent to the port. If data is not yet available, one event is sent to the port when data first becomes available.</p> <p>When an event for a <code>PORT_SOURCE_FD</code> object is retrieved, the object no longer has an association with the port. The event can be processed without the possibility that another thread can retrieve a subsequent event for the same object. After processing of the file descriptor is completed, the <code>port_associate()</code> function can be called to reassociate the object with the port.</p> <p>The parent and child processes are allowed to retrieve events from file descriptors shared after a call to <code>fork(2)</code>. The process performing the first association with a port (parent or child process) is designated as the owner of the association. Only the owner of an association is allowed to dissociate the file descriptor from a port. The association is removed if the owner of the association closes the port.</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 <code>port_associate()</code> and <code>port_dissociate()</code> functions will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>EBADF</code></td> <td>The <i>port</i> identifier is not valid.</td> </tr> <tr> <td><code>EBADFD</code></td> <td>The <i>source</i> argument is of type <code>PORT_SOURCE_FD</code> and the object argument is not a valid file descriptor.</td> </tr> <tr> <td><code>EINVAL</code></td> <td>The <i>source</i> argument is not valid.</td> </tr> </table>	<code>EBADF</code>	The <i>port</i> identifier is not valid.	<code>EBADFD</code>	The <i>source</i> argument is of type <code>PORT_SOURCE_FD</code> and the object argument is not a valid file descriptor.	<code>EINVAL</code>	The <i>source</i> argument is not valid.
<code>EBADF</code>	The <i>port</i> identifier is not valid.						
<code>EBADFD</code>	The <i>source</i> argument is of type <code>PORT_SOURCE_FD</code> and the object argument is not a valid file descriptor.						
<code>EINVAL</code>	The <i>source</i> argument is not valid.						

port_associate(3C)

The `port_associate()` function will fail if:

EAGAIN The maximum number of objects associated with the port was exceeded. The maximum allowable number of events or association of objects per port is the minimum value of the `process.max-port-events` resource control at the time `port_create(3C)` was used to create the port. See `setrctl(2)` and `rctladm(1M)` for information on using resource controls.

The number of objects associated with a port is composed of all supported resource types. Some of the source types do not explicitly use the `port_associate()` function.

ENOMEM The physical memory limits of the system have been exceeded.

The `port_dissociate()` function will fail if:

EACCES The process is not the owner of the association.

EXAMPLES

EXAMPLE 1 Retrieve data from a pipe file descriptor.

The following example retrieves data from a pipe file descriptor.

```
#include <port.h>

int          port;
int          fd;
int          error;
int          index;
void         *mypointer;
port_event_t pev;
struct timespec_t timeout;
char         rbuf[STRSIZE];
int          fds[MAXINDEX];

/* create a port */
port = port_create();

for (index = 0; index < MAXINDEX; index++) {
    error = mkfifo(name[index], S_IRWXU | S_IRWXG | S_IRWXO);
    if (error)
        /* handle error code */
        fds[index] = open(name[index], O_RDWR);

    /* associate pipe file descriptor with the port */
    error = port_associate(port, PORT_SOURCE_FD, fds[index],
        POLLIN, mypointer);
}
...
timeout.tv_sec = 1;      /* user defined */
timeout.tv_nsec = 0;

/* loop to retrieve data from the list of pipe file descriptors */
for (...) {
    /* retrieve a single event */
```

EXAMPLE 1 Retrieve data from a pipe file descriptor. (Continued)

```

error = port_get(port, &pev, &timeout);
if (error) {
    /* handle error code */
}
fd = pev.portev_object;
if (read(fd, rbuf, STRSIZE)) {
    /* handle error code */
}
if (fd-still-accepting-data) {
    /*
     * re-associate the file descriptor with the port.
     * The re-association is required for the
     * re-activation of the data detection.
     * Internals events and user arguments are set to the
     * new (or the same) values delivered here.
     */
    error = port_associate(port, PORT_SOURCE_FD, fd, POLLIN,
        pev.portev_user);
} else {
    /*
     * If file descriptor is no longer required,
     * - it can remain disabled but still associated with
     *   the port, or
     * - it can be dissociated from the port.
     */
}

```

EXAMPLE 2 Bind AIO transaction to a specific port.

The following example binds the AIO transaction to a specific port.

```

#include <port.h>

int          port;
port_notify_t pn;
aiocb_t     aiocb;
aiocb_t     *aiocbp;
void        *mypointer;
int         error;
int         my_errno;
int         my_status;
struct timespec timeout;
port_event_t pev;

port = port_create();
...
/* fill AIO specific part */
aiocb.aio_fildes = fd;
aiocb.aio_nbytes = BUFSIZE;
aiocb.aio_buf = bufp;
aiocb.aio_offset = 0;
aiocb.aio_sigevent.sigev_notify = SIGEV_NONE; /* SIGV_SIGNAL, ... */
aiocb.aio_sigevent.sigev_signo = <signal-no>; /* if SIGV_SIGNAL */

```

port_associate(3C)

EXAMPLE 2 Bind AIO transaction to a specific port. (Continued)

```
/* port specific part */
pn.portnfy_port = port;
pn.portnfy_user = mypointer;

aiocb.aio_sigevent.sigev_notify = SIGEV_PORT;
aiocb.aio_sigevent.sigev_value.sival_ptr = &pn

/*
 * The aio_read() function binds internally the asynchronous I/O
 * transaction with the port delivered in port_notify_t.
 */
error = aio_read(&aiocb);

timeout.tv_sec = 1;      /* user defined */
timeout.tv_nsec = 0;

/* retrieve a single event */
error = port_get(port, &pev, &timeout);
if (error) {
    /* handle error code */
}

/*
 * pev.portev_object contains a pointer to the aiocb structure
 * delivered in port_notify_t (see aio_read()).
 */
aiocbp = pev.portev_object;

/* check error code and return value in
my_errno = aio_error(aiocbp);
...
my_status = aio_return(aiocbp);
...

```

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	all
Availability	SUNWcsr, SUNWhea
Interface Stability	Evolving
MT-Level	Safe

SEE ALSO [rctladm\(1M\)](#), [poll\(2\)](#), [setrctl\(2\)](#), [port_alert\(3C\)](#), [port_create\(3C\)](#), [port_get\(3C\)](#), [port_send\(3C\)](#), [attributes\(5\)](#)

NAME	port_create – create a port																		
SYNOPSIS	<pre>#include <port.h> int port_create(void);</pre>																		
DESCRIPTION	<p>The <code>port_create()</code> function establishes a queue that multiplexes events from disjoint sources. Each source has a corresponding object type and source-specific mechanism for associating an object with a port.</p> <table border="1"> <thead> <tr> <th>source</th> <th>object type</th> <th>association mechanism</th> </tr> </thead> <tbody> <tr> <td>PORT_SOURCE_AIO</td> <td>struct aiocb</td> <td>aio_read(3RT), aio_write(3RT), lio_listio(3RT)</td> </tr> <tr> <td>PORT_SOURCE_FD</td> <td>file descriptor</td> <td>port_associate(3C)</td> </tr> <tr> <td>PORT_SOURCE_TIMER</td> <td>timer_t</td> <td>timer_create(3RT)</td> </tr> <tr> <td>PORT_SOURCE_USER</td> <td>uintptr_t</td> <td>port_send(3C)</td> </tr> <tr> <td>PORT_SOURCE_ALERT</td> <td>uintptr_t</td> <td>port_alert(3C)</td> </tr> </tbody> </table> <p>PORT_SOURCE_AIO events represent the completion of an asynchronous I/O transaction. An asynchronous I/O transaction is associated with a port by specifying SIGEV_PORT as its notification mechanism. See <code>aio_read(3RT)</code>, <code>aio_write(3RT)</code>, <code>lio_listio(3RT)</code>, and <code>aio.h(3HEAD)</code> for details.</p> <p>PORT_SOURCE_FD events represent a transition in the <code>poll(2)</code> status of a given file descriptor. Once an event is delivered, the file descriptor is no longer associated with the port. A file descriptor is associated (or re-associated) with a port using the <code>port_associate(3C)</code> function.</p> <p>PORT_SOURCE_TIMER events represent one or more timer expirations for a given timer. A timer is associated with a port by specifying SIGEV_PORT as its notification mechanism. See <code>timer_create(3RT)</code> for more information.</p> <p>PORT_SOURCE_USER events represent user-defined events. These events are generated by <code>port_send(3C)</code> or <code>port_sendn(3C)</code>.</p> <p>PORT_SOURCE_ALERT events indicate that the port itself is in alert mode. The mode of the port is changed with <code>port_alert(3C)</code>. The <code>port_create()</code> function returns a file descriptor that represents a newly created port. The <code>close(2)</code> function destroys the port and frees all allocated resources.</p> <p>The <code>port_get(3C)</code> and <code>port_getn(3C)</code> functions retrieve events from a port. They ignore non retrievable events (non-own or non-shareable events).</p>	source	object type	association mechanism	PORT_SOURCE_AIO	struct aiocb	aio_read(3RT), aio_write(3RT), lio_listio(3RT)	PORT_SOURCE_FD	file descriptor	port_associate(3C)	PORT_SOURCE_TIMER	timer_t	timer_create(3RT)	PORT_SOURCE_USER	uintptr_t	port_send(3C)	PORT_SOURCE_ALERT	uintptr_t	port_alert(3C)
source	object type	association mechanism																	
PORT_SOURCE_AIO	struct aiocb	aio_read(3RT), aio_write(3RT), lio_listio(3RT)																	
PORT_SOURCE_FD	file descriptor	port_associate(3C)																	
PORT_SOURCE_TIMER	timer_t	timer_create(3RT)																	
PORT_SOURCE_USER	uintptr_t	port_send(3C)																	
PORT_SOURCE_ALERT	uintptr_t	port_alert(3C)																	

port_create(3C)

As a port is represented by a file descriptor, ports are shared between child and parent processes after `fork()`. Both can continue to associate sources with the port, both can receive events from the port, and events associated with and/or generated by either process are retrievable in the other. Since some events might not have meaning in both parent and child, care must be taken when using ports after `fork()`.

If a port is exported to other processes, the port is destroyed on last close.

`PORT_SOURCE_USER` and `PORT_SOURCE_ALERT` events can be distributed across processes. `PORT_SOURCE_FD` events can only be shared between processes when child processes inherit opened file descriptors from the parent process. See `fork(2)`. `PORT_SOURCE_TIMER` and `PORT_SOURCE_AIO` cannot be shared between processes.

RETURN VALUES

Upon successful completion, the `port_create()` function returns a non-negative value, the port identifier. Otherwise, `-1` is returned and `errno` is set to indicate the error.

ERRORS

The `port_create()` function will fail if:

EAGAIN The maximum allowable number of ports is currently open in the system. The maximum allowable number of ports is the minimum value of the `project.max-port-ids` resource control. See `setrctl(2)` and `rctladm(1M)` for information on using resource controls.

EMFILE The process has too many open descriptors.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	all
Availability	SUNWcsr, SUNWhea
Interface Stability	Evolving
MT-Level	Safe

SEE ALSO

`rctladm(1M)`, `close(2)`, `exit(2)`, `fork(2)`, `poll(2)`, `setrctl(2)`, `aio_read(3RT)`, `aio_write(3RT)`, `aio.h(3HEAD)`, `lio_listio(3RT)`, `timer_create(3RT)`, `attributes(5)`

NAME	port_get, port_getn – retrieve event information from a port
SYNOPSIS	<pre>#include <port.h> int port_get(int port, port_event_t *pe, const timespec_t *timeout); int port_getn(int port, port_event_t list[], uint_t max, uint_t *nget, const timespec_t *timeout);</pre>
DESCRIPTION	<p>The port_get() and port_getn() functions retrieve events from a port. The port_get() function retrieves at most a single event. The port_getn() function can retrieve multiple events.</p> <p>The <i>pe</i> argument points to an uninitialized port_event_t structure that is filled in by the system when the port_get() function returns successfully.</p> <p>The port_event_t structure contains the following members:</p> <pre>int portev_events; /* detected events */ ushort_t portev_source; /* event source */ uintptr_t portev_object; /* specific to event source */ void *portev_user; /* user defined cookie */</pre> <p>The portev_events and portev_object members are specific to the event source. The portev_events denotes the delivered events. The portev_object refers to the associated object (see port_create(3C)). The portev_source member specifies the source of the event. The portev_user member is a user-specified value.</p> <p>If the <i>timeout</i> pointer is NULL, the port_get() function blocks until an event is available. To poll for an event without waiting, <i>timeout</i> should point to a zeroed timespec. A non-zeroed timespec specifies the desired time to wait for events. The port_get() function returns before the timeout elapses if an event is available, a signal occurs, a port is closed by another thread, or the port is in or enters alert mode. See port_alert(3C) for details on alert mode.</p> <p>The port_getn() function can retrieve multiple events from a port. The <i>list</i> argument is an array of uninitialized port_event_t structures that is filled in by the system when the port_getn() function returns successfully. The <i>nget</i> argument points to the desired number of events to be retrieved. The <i>max</i> parameter specifies the maximum number of events that can be returned in <i>list</i>[]. If <i>max</i> is 0, the value pointed to by <i>nget</i> is set to the number of events available on the port. The port_getn() function returns immediately but no events are retrieved.</p> <p>The port_getn() function block until the desired number of events are available, the timeout elapses, a signal occurs, a port is closed by another thread, or the port is in or enters alert mode.</p> <p>On return, the value pointed to by <i>nget</i> is updated to the actual number of events retrieved in <i>list</i>.</p>

port_get(3C)

Threads calling the `port_get()` function might starve threads waiting in the `port_getn()` function for more than one event. Similarly, threads calling the `port_getn()` function for n events might starve threads waiting in the `port_getn()` function for more than n events.

The `port_get()` and the `port_getn()` functions ignore non-shareable events (see [port_create\(3C\)](#)) generated by other processes.

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

ERRORS The `port_get()` and `port_getn()` functions will fail if:

EBADF	The <i>port</i> identifier is not valid.
EBADFD	The <i>port</i> argument is not an event port file descriptor.
EFAULT	Event or event list can not be delivered (<i>list[]</i> pointer and/or user space reserved to accommodate the list of events is not reasonable), or the <i>timeout</i> argument is not reasonable.
EINTR	A signal was caught during the execution of the function.
EINVAL	The <i>timeout</i> element <i>tv_sec</i> is < 0 or the <i>timeout</i> element <i>tv_nsec</i> is < 0 or > 1000000000.
ETIME	The time interval expired before the expected number of events have been posted to the port.

The `port_getn()` function will fail if:

EINVAL	The <i>list[]</i> argument is NULL, the <i>nget</i> argument is NULL, or the content of <i>nget</i> is > <i>max</i> and <i>max</i> is > 0.
EFAULT	The <i>timeout</i> argument is not reasonable.
ETIME	The time interval expired before the expected number of events have been posted to the port (original value in <i>nget</i>), or <i>nget</i> is updated with the number of returned <code>port_event_t</code> structures in <i>list[]</i> .

EXAMPLES **EXAMPLE 1** Send a user event (PORT_SOURCE_USER) to a port and retrieve it with `port_get()`.

The following example sends a user event (PORT_SOURCE_USER) to a port and retrieves it with `port_get()`. The `portev_user` and `portev_events` members of the `port_event_t` structure are the same as the corresponding user and events arguments of the [port_send\(3C\)](#) function.

```
#include <port.h>

int          myport;
port_event_t pe;
struct timespec timeout;
int          ret;
```

EXAMPLE 1 Send a user event (PORT_SOURCE_USER) to a port and retrieve it with port_get(). (Continued)

```

void          *user;
uintptr_t     object;

myport = port_create();
if (myport < 0) {
    /* port creation failed ... */
    ...
    return(...);
}
...
events = 0x01;          /* own event definition(s) */
object = <myobject>;
user = <my_own_value>;
ret = port_send(myport, events, object, user);
if (ret == -1) {
    /* error detected ... */
    ...
    close(myport);
    return (...);
}

/*
 * The following code could also be executed in another thread or
 * process.
 */
timeout.tv_sec = 1;     /* user defined */
timeout.tv_nsec = 0;
ret = port_get(myport, &pe, &timeout);
if (ret == -1) {
    /*
     * error detected :
     * - EINTR or ETIME : log error code and try again ...
     * - Other kind of errors : may have to close the port ...
     */
    return(...);
}

/*
 * After port_get() returns successfully, the port_event_t
 * structure will be filled with:
 * pe.portev_source = PORT_SOURCE_USER
 * pe.portev_events = 0x01
 * pe.portev_object = <myobject>
 * pe.portev_user = <my_own_value>
 */
...
close(myport);

```

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

port_get(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	all
Availability	SUNWcsr, SUNWhea
Interface Stability	Evolving
MT-Level	Safe

SEE ALSO [port_alert\(3C\)](#), [port_associate\(3C\)](#), [port_create\(3C\)](#), [port_send\(3C\)](#), [attributes\(5\)](#)

NAME	port_send, port_sendn – send a user-defined event to a port or list of ports										
SYNOPSIS	<pre>#include <port.h> int port_send(int port, int events, void *user); int port_sendn(int ports[], int errors[], uint_t nent, int events, void *user);</pre>										
DESCRIPTION	<p>The port_send() function submits a user-defined event to a specified port. The <i>port</i> argument is a file descriptor that represents a port. The sent event has its portev_events member set to the value specified in the <i>events</i> parameter and its portev_user member set to the value specified in the <i>user</i> parameter. The portev_object member of an event sent with port_send() is unspecified.</p> <p>The port_sendn() function submits a user-defined event to multiple ports. The <i>ports</i> argument is an array of file descriptors that represents ports (see port_create(3C)). The <i>nent</i> argument specifies the number of file descriptors in the <i>ports</i>[] array. An event is submitted to each specified port. Each event has its portev_events member set to the value specified in the <i>events</i> parameter and its portev_user member set to the value specified in the <i>user</i> parameter. The portev_object member of <i>events</i> sent with port_sendn() is unspecified.</p> <p>A port that is in alert mode can be sent an event, but that event will not be retrievable until the port has resumed normal operation. See port_alert(3C).</p>										
RETURN VALUES	<p>Upon successful completion, the port_send() function returns 0. Otherwise, it returns -1 and sets errno to indicate the error.</p> <p>The port_sendn() function returns the number of successfully submitted events. A non-negative return value less than the <i>nent</i> argument indicates that at least one error occurred. In this case, each element of the <i>errors</i>[] array is filled in. An element of the <i>errors</i>[] array is set to 0 if the event was successfully sent to the corresponding port in the <i>ports</i>[] array, or is set to indicate the error if the event was not successfully sent. If an error occurs, the port_sendn() function returns -1 and sets errno to indicate the error.</p>										
ERRORS	<p>The port_send() and port_sendn() functions will fail if:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">EAGAIN</td> <td>The maximum number of events per port is exceeded. The maximum allowable number of events per port is the minimum value of the process.max-port-events resource control at the time port_create(3C) was used to create the port.</td> </tr> <tr> <td>EBADF</td> <td>The port file descriptor is not valid.</td> </tr> <tr> <td>EBADFD</td> <td>The <i>port</i> argument is not an event port file descriptor.</td> </tr> <tr> <td>ENOMEM</td> <td>There is not enough memory available to satisfy the request.</td> </tr> </table> <p>The port_sendn() function will fail if:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">EFAULT</td> <td>The <i>ports</i>[] pointer or <i>errors</i>[] pointer is not reasonable.</td> </tr> </table>	EAGAIN	The maximum number of events per port is exceeded. The maximum allowable number of events per port is the minimum value of the process.max-port-events resource control at the time port_create(3C) was used to create the port.	EBADF	The port file descriptor is not valid.	EBADFD	The <i>port</i> argument is not an event port file descriptor.	ENOMEM	There is not enough memory available to satisfy the request.	EFAULT	The <i>ports</i> [] pointer or <i>errors</i> [] pointer is not reasonable.
EAGAIN	The maximum number of events per port is exceeded. The maximum allowable number of events per port is the minimum value of the process.max-port-events resource control at the time port_create(3C) was used to create the port.										
EBADF	The port file descriptor is not valid.										
EBADFD	The <i>port</i> argument is not an event port file descriptor.										
ENOMEM	There is not enough memory available to satisfy the request.										
EFAULT	The <i>ports</i> [] pointer or <i>errors</i> [] pointer is not reasonable.										

port_send(3C)

EINVAL The value of the *nent* argument is 0.

EXAMPLES

EXAMPLE 1 Use `port_send()` to send a user event (`PORT_SOURCE_USER`) to a port.

The following example uses `port_send()` to send a user event (`PORT_SOURCE_USER`) to a port and `port_get()` to retrieve it. The `portev_user` and `portev_events` members of the `port_event_t` structure are the same as the corresponding `user` and `events` arguments of the `port_send()` function.

```
#include <port.h>

int          myport;
port_event_t pe;
struct timespec timeout;
int          ret;
void         *user;

myport = port_create();
if (myport) {
    /* port creation failed ... */
    ...
    return(...);
}
...
events = 0x01;          /* own event definition(s) */
user = <my_own_value>;
ret = port_send(myport, events, user);
if (ret == -1) {
    /* error detected ... */
    ...
    close(myport);
    return (...);
}

/*
 * The following code could also be executed from another thread or
 * process.
 */
timeout.tv_sec = 1;    /* user defined */
timeout.tv_nsec = 0;
ret = port_get(myport, &pe, &timeout);
if (ret == -1) {
    /*
     * error detected :
     * - EINTR or ETIME : log error code and try again ...
     * - Other kind of errors : may have to close the port ...
     */
    return(...);
}

/*
 * After port_get() returns successfully, the port_event_t
 * structure will be filled with:
 * pe.portev_source = PORT_SOURCE_USER
 * pe.portev_events = 0x01
 * pe.portev_object = unspecified
 */
```

port_send(3C)

EXAMPLE 1 Use port_send() to send a user event (PORT_SOURCE_USER) to a port.
(Continued)

```
* pe.portev_user = <my_own_value>
*/
...
close(myport);
```

USAGE See setrctl(2) and rctladm(1M) for information on using resource controls.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	all
Availability	SUNWcsr, SUNWhea
Interface Stability	Evolving
MT-Level	Safe

SEE ALSO rctladm(1M), setrctl(2), port_alert(3C), port_associate(3C), port_create(3C), port_get(3C), attributes(5)

posix_openpt(3C)

NAME	posix_openpt – open a pseudo terminal device										
SYNOPSIS	<pre>#include <stdlib.h> #include <fcntl.h> int posix_openpt(int <i>oflag</i>);</pre>										
DESCRIPTION	<p>The <code>posix_openpt()</code> function establishes a connection between a master device for a pseudo-terminal and a file descriptor. The file descriptor is used by other I/O functions that refer to that pseudo-terminal.</p> <p>The file status flags and file access modes of the open file description are 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>.</p> <table><tr><td><code>O_RDWR</code></td><td>Open for reading and writing.</td></tr><tr><td><code>O_NOCTTY</code></td><td>If set, <code>posix_openpt()</code> does not cause the terminal device to become the controlling terminal for the process.</td></tr></table> <p>The behavior of other values for the <i>oflag</i> argument is unspecified.</p>	<code>O_RDWR</code>	Open for reading and writing.	<code>O_NOCTTY</code>	If set, <code>posix_openpt()</code> does not cause the terminal device to become the controlling terminal for the process.						
<code>O_RDWR</code>	Open for reading and writing.										
<code>O_NOCTTY</code>	If set, <code>posix_openpt()</code> does not cause the terminal device to become the controlling terminal for the process.										
RETURN VALUES	Upon successful completion, the <code>posix_openpt()</code> function opens a master pseudo-terminal device and returns a non-negative integer representing the lowest numbered unused file descriptor. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.										
ERRORS	<p>The <code>posix_openpt()</code> function will fail if:</p> <table><tr><td><code>EMFILE</code></td><td>{<code>OPEN_MAX</code>} file descriptors are currently open in the calling process.</td></tr><tr><td><code>ENFILE</code></td><td>The maximum allowable number of files is currently open in the system.</td></tr></table> <p>The <code>posix_openpt()</code> function may fail if:</p> <table><tr><td><code>EINVAL</code></td><td>The value of <i>oflag</i> is not valid.</td></tr><tr><td><code>EAGAIN</code></td><td>Out of pseudo-terminal resources.</td></tr><tr><td><code>ENOSR</code></td><td>Out of STREAMS resources.</td></tr></table>	<code>EMFILE</code>	{ <code>OPEN_MAX</code> } file descriptors are currently open in the calling process.	<code>ENFILE</code>	The maximum allowable number of files is currently open in the system.	<code>EINVAL</code>	The value of <i>oflag</i> is not valid.	<code>EAGAIN</code>	Out of pseudo-terminal resources.	<code>ENOSR</code>	Out of STREAMS resources.
<code>EMFILE</code>	{ <code>OPEN_MAX</code> } file descriptors are currently open in the calling process.										
<code>ENFILE</code>	The maximum allowable number of files is currently open in the system.										
<code>EINVAL</code>	The value of <i>oflag</i> is not valid.										
<code>EAGAIN</code>	Out of pseudo-terminal resources.										
<code>ENOSR</code>	Out of STREAMS resources.										
EXAMPLES	<p>EXAMPLE 1 Open a pseudo-terminal.</p> <p>The following example opens a pseudo-terminal and returns the name of the slave device and a file descriptor.</p> <pre>#include fcntl.h> #include stdio.h> int masterfd, slavefd; char *slavedevice;</pre>										

EXAMPLE 1 Open a pseudo-terminal. (Continued)

```

masterfd = posix_openpt(O_RDWR|O_NOCTTY);

if (masterfd == -1
    || grantpt (masterfd) == -1
    || unlockpt (masterfd) == -1
    || (slavedevice = ptsname (masterfd)) == NULL)
    return -1;

printf("slave device is: %s\n", slavedevice);

slavefd = open(slave, O_RDWR|O_NOCTTY);
if (slavefd < 0)
    return -1;

```

USAGE This function provides a method for portably obtaining a file descriptor of a master terminal device for a pseudo-terminal. The [grantpt\(3C\)](#) and [ptsname\(3C\)](#) functions can be used to manipulate mode and ownership permissions and to obtain the name of the slave device, respectively.

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [open\(2\)](#), [grantpt\(3C\)](#), [ptsname\(3C\)](#), [unlockpt\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

posix_spawn(3C)

NAME	posix_spawn, posix_spawnnp – spawn a process
SYNOPSIS	<pre>#include <spawn.h> int posix_spawn(pid_t *restrict <i>pid</i>, const char *restrict <i>path</i>, const posix_spawn_file_actions_t *<i>file_actions</i>, const posix_spawnattr_t *restrict <i>attrp</i>, char *const <i>argv</i>[restrict], char *const <i>envp</i>[restrict]); int posix_spawnnp(pid_t *restrict <i>pid</i>, const char *restrict <i>file</i>, const posix_spawn_file_actions_t *<i>file_actions</i>, const posix_spawnattr_t *restrict <i>attrp</i>, char *const <i>argv</i>[restrict], char *const <i>envp</i>[restrict]);</pre>
DESCRIPTION	<p>The <code>posix_spawn()</code> and <code>posix_spawnnp()</code> functions create a new process (child process) from the specified process image. The new process image is constructed from a regular executable file called the new process image file.</p> <p>When a C program is executed as the result of this call, it is entered as a C language function call as follows:</p> <pre>int main(int <i>argc</i>, char *<i>argv</i>[]);</pre> <p>where <i>argc</i> is the argument count and <i>argv</i> is an array of character pointers to the arguments themselves. In addition, the following variable</p> <pre>extern char **<i>environ</i>;</pre> <p>is initialized as a pointer to an array of character pointers to the environment strings.</p> <p>The argument <i>argv</i> is an array of character pointers to null-terminated strings. The last member of this array is a null pointer and is not counted in <i>argc</i>. These strings constitute the argument list available to the new process image. The value in <i>argv</i>[0] should point to a filename that is associated with the process image being started by the <code>posix_spawn()</code> or <code>posix_spawnnp()</code> function.</p> <p>The argument <i>envp</i> is an array of character pointers to null-terminated strings. These strings constitute the environment for the new process image. The environment array is terminated by a null pointer.</p> <p>The number of bytes available for the child process's combined argument and environment lists is {ARG_MAX}, counting all character pointers, the strings they point to, the trailing null bytes in the strings, and the list-terminating null pointers. There is no additional system overhead included in this total.</p> <p>The <i>path</i> argument to <code>posix_spawn()</code> is a pathname that identifies the new process image file to execute.</p>

The *file* parameter to `posix_spawn()` is used to construct a pathname that identifies the new process image file. If the *file* parameter contains a slash character, the *file* parameter is used as the pathname for the new process image file. Otherwise, the path prefix for this file is obtained by a search of the directories passed as the environment variable `PATH`. If this environment variable is not defined, the results of the search are implementation-defined.

If *file_actions* is a null pointer, then file descriptors open in the calling process remain open in the child process, except for those whose close-on-exec flag `FD_CLOEXEC` is set (see `fcntl(2)`). For those file descriptors that remain open, all attributes of the corresponding open file descriptions, including file locks (see `fcntl(2)`), remain unchanged.

If *file_actions* is not `NULL`, then the file descriptors open in the child process are those open in the calling process as modified by the spawn file actions object pointed to by *file_actions* and the `FD_CLOEXEC` flag of each remaining open file descriptor after the spawn file actions have been processed. The effective order of processing the spawn file actions are:

1. The set of open file descriptors for the child process are initially the same set as is open for the calling process. All attributes of the corresponding open file descriptions, including file locks (see `fcntl(2)`), remain unchanged.
2. The signal mask, signal default actions, and the effective user and group IDs for the child process are changed as specified in the attributes object referenced by *attrp*.
3. The file actions specified by the spawn file actions object are performed in the order in which they were added to the spawn file actions object.
4. Any file descriptor that has its `FD_CLOEXEC` flag set (see `fcntl(2)`) is closed.

The `posix_spawnattr_t` spawn attributes object type is defined in `<spawn.h>`. It contains at least the attributes defined below.

If the `POSIX_SPAWN_SETPGROUP` flag is set in the *spawn-flags* attribute of the object referenced by *attrp*, and the *spawn-pgroup* attribute of the same object is non-zero, then the child's process group is as specified in the *spawn-pgroup* attribute of the object referenced by *attrp*.

As a special case, if the `POSIX_SPAWN_SETPGROUP` flag is set in the *spawn-flags* attribute of the object referenced by *attrp*, and the *spawn-pgroup* attribute of the same object is set to zero, then the child will be in a new process group with a process group ID equal to its process ID.

If the `POSIX_SPAWN_SETPGROUP` flag is not set in the *spawn-flags* attribute of the object referenced by *attrp*, the new child process inherits the parent's process group.

If the `POSIX_SPAWN_SETSCHEDPARAM` flag is set in the *spawn-flags* attribute of the object referenced by *attrp*, but `POSIX_SPAWN_SETSCHEDULER` is not set, the new process image initially has the scheduling policy of the calling process with the scheduling parameters specified in the *spawn-schedparam* attribute of the object referenced by *attrp*.

posix_spawn(3C)

If the `POSIX_SPAWN_SETSCHEDULER` flag is set in *spawn-flags* attribute of the object referenced by *attrp* (regardless of the setting of the `POSIX_SPAWN_SETSCHEDPARAM` flag), the new process image initially has the scheduling policy specified in the *spawn-schedpolicy* attribute of the object referenced by *attrp* and the scheduling parameters specified in the *spawn-schedparam* attribute of the same object.

The `POSIX_SPAWN_RESETEUIDS` flag in the *spawn-flags* attribute of the object referenced by *attrp* governs the effective user ID of the child process. If this flag is not set, the child process inherits the parent process's effective user ID. If this flag is set, the child process's effective user ID is reset to the parent's real user ID. In either case, if the set-user-ID mode bit of the new process image file is set, the effective user ID of the child process becomes that file's owner ID before the new process image begins execution. If this flag is set, the child process's effective user ID is reset to the parent's real user ID. In either case, if the set-user-ID mode bit of the new process image file is set, the effective user ID of the child process becomes that file's owner ID before the new process image begins execution.

The `POSIX_SPAWN_RESETEGIDS` flag in the *spawn-flags* attribute of the object referenced by *attrp* also governs the effective group ID of the child process. If this flag is not set, the child process inherits the parent process's effective group ID. If this flag is set, the child process's effective group ID is reset to the parent's real group ID. In either case, if the set-group-ID mode bit of the new process image file is set, the effective group ID of the child process becomes that file's group ID before the new process image begins execution.

If the `POSIX_SPAWN_SETSIGMASK` flag is set in the *spawn-flags* attribute of the object referenced by *attrp*, the child process initially has the signal mask specified in the *spawn-sigmask* attribute of the object referenced by *attrp*.

If the `POSIX_SPAWN_SETSIGDEF` flag is set in the *spawn-flags* attribute of the object referenced by *attrp*, the signals specified in the *spawn-sigdefault* attribute of the same object is set to their default actions in the child process. Signals set to the default action in the parent process are set to the default action in the child process.

Signals set to be caught by the calling process are set to the default action in the child process.

Except for `SIGCHLD`, signals set to be ignored by the calling process image are set to be ignored by the child process, unless otherwise specified by the `POSIX_SPAWN_SETSIGDEF` flag being set in the *spawn-flags* attribute of the object referenced by *attrp* and the signals being indicated in the *spawn-sigdefault* attribute of the object referenced by *attrp*.

If the `SIGCHLD` signal is set to be ignored by the calling process, it is unspecified whether the `SIGCHLD` signal is set to be ignored or to the default action in the child process, unless otherwise specified by the `POSIX_SPAWN_SETSIGDEF` flag being set in the *spawn-flags* attribute of the object referenced by *attrp* and the `SIGCHLD` signal being indicated in the *spawn-sigdefault* attribute of the object referenced by *attrp*.

If the value of the *attrp* pointer is `NULL`, then the default values are used.

All process attributes, other than those influenced by the attributes set in the object referenced by *attrp* as specified above or by the file descriptor manipulations specified in *file_actions* appear in the new process image as though `fork(2)` had been called to create a child process and then a member of the `exec(2)` family of functions had been called by the child process to execute the new process image.

The fork handlers are not run when `posix_spawn()` or `posix_spawnp()` is called.

RETURN VALUES

Upon successful completion, `posix_spawn()` and `posix_spawnp()` return the process ID of the child process to the parent process in the variable pointed to by a non-null *pid* argument, and return zero as the function return value. Otherwise, no child process is created, the value stored into the variable pointed to by a non-null *pid* is unspecified, and an error number is returned as the function return value to indicate the error. If the *pid* argument is a null pointer, the process ID of the child is not returned to the caller.

ERRORS

The `posix_spawn()` and `posix_spawnp()` functions will fail if:

EINVAL The value specified by *file_actions* or *attrp* is invalid.

If this error occurs after the calling process successfully returns from the `posix_spawn()` or `posix_spawnp()` function, the child process might exit with exit status 127.

If `posix_spawn()` or `posix_spawnp()` fails for any of the reasons that would cause `fork()` or one of the `exec` family of functions to fail, an error value is returned as described by `fork()` and `exec`, respectively (or, if the error occurs after the calling process successfully returns, the child process exits with exit status 127).

If `POSIX_SPAWN_SETPGROUP` is set in the *spawn-flags* attribute of the object referenced by *attrp*, and `posix_spawn()` or `posix_spawnp()` fails while changing the child's process group, an error value is returned as described by `setpgid(2)` (or, if the error occurs after the calling process successfully returns, the child process exits with exit status 127).

If `POSIX_SPAWN_SETSCHEDPARAM` is set and `POSIX_SPAWN_SETSCHEDULER` is not set in the *spawn-flags* attribute of the object referenced by *attrp*, then if `posix_spawn()` or `posix_spawnp()` fails for any of the reasons that would cause `sched_setparam(3RT)` to fail, an error value is returned as described by `sched_setparam()` (or, if the error occurs after the calling process successfully returns, the child process exits with exit status 127).

If `POSIX_SPAWN_SETSCHEDULER` is set in the *spawn-flags* attribute of the object referenced by *attrp*, and if `posix_spawn()` or `posix_spawnp()` fails for any of the reasons that would cause `sched_setscheduler(3RT)` to fail, an error value is returned as described by `sched_setscheduler()` (or, if the error occurs after the calling process successfully returns, the child process exits with exit status 127).

posix_spawn(3C)

If the *file_actions* argument is not NULL and specifies any `close(2)`, `dup2(3C)`, or `open(2)` actions to be performed, and if `posix_spawn()` or `posix_spawnnp()` fails for any of the reasons that would cause `close()`, `dup2()`, or `open()` to fail, an error value is returned as described by `close()`, `dup2()`, and `open()`, respectively (or, if the error occurs after the calling process successfully returns, the child process exits with exit status 127). An open file action might, by itself, result in any of the errors described by `close()` or `dup2()`, in addition to those described by `open()`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `alarm(2)`, `chmod(2)`, `close(2)`, `dup(2)`, `exec(2)`, `exit(2)`, `fcntl(2)`, `fork(2)`, `kill(2)`, `open(2)`, `setpgid(2)`, `setuid(2)`, `stat(2)`, `times(2)`, `dup2(3C)`, `posix_spawn_file_actions_addclose(3C)`, `posix_spawn_file_actions_adddup2(3C)`, `posix_spawn_file_actions_addopen(3C)`, `posix_spawn_file_actions_destroy(3C)`, `posix_spawn_file_actions_init(3C)`, `posix_spawnattr_destroy(3C)`, `posix_spawnattr_init(3C)`, `posix_spawnattr_getsigdefault(3C)`, `posix_spawnattr_getflags(3C)`, `posix_spawnattr_getpgroup(3C)`, `posix_spawnattr_getschedparam(3C)`, `posix_spawnattr_getschedpolicy(3C)`, `posix_spawnattr_getsigmask(3C)`, `posix_spawnattr_setflags(3C)`, `posix_spawnattr_setpgroup(3C)`, `posix_spawnattr_setschedparam(3C)`, `posix_spawnattr_setschedpolicy(3C)`, `posix_spawnattr_setsigdefault(3C)`, `posix_spawnattr_setsigmask(3C)`, `sched_setparam(3RT)`, `sched_setscheduler(3RT)`, `wait(3C)`, `attributes(5)`, `standards(5)`

NAME posix_spawnattr_destroy, posix_spawnattr_init – destroy and initialize spawn attributes object

SYNOPSIS

```
#include <spawn.h>

int posix_spawnattr_destroy(posix_spawnattr_t *attr);
int posix_spawnattr_init(posix_spawnattr_t *attr);
```

DESCRIPTION The `posix_spawnattr_destroy()` function destroys a spawn attributes object. A destroyed `attr` attributes object can be reinitialized using `posix_spawnattr_init()`. The results of otherwise referencing the object after it has been destroyed are undefined. An implementation can cause `posix_spawnattr_destroy()` to set the object referenced by `attr` to an invalid value.

The `posix_spawnattr_init()` function initializes a spawn attributes object `attr` with the default value for all of the individual attributes used by the implementation. Results are undefined if `posix_spawnattr_init()` is called specifying an already initialized `attr` attributes object.

A spawn attributes object is of type `posix_spawnattr_t` (defined in `<spawn.h>`) and is used to specify the inheritance of process attributes across a spawn operation.

No attributes other than those defined by IEEE Std 1003.1-200x are provided.

The resulting spawn attributes object (possibly modified by setting individual attribute values), is used to modify the behavior of `posix_spawn(3C)` or `posix_spawnp(3C)`. After a spawn attributes object has been used to spawn a process by a call to `posix_spawn()` or `posix_spawnp()`, any function affecting the attributes object (including destruction) will not affect any process that has been spawned in this way.

RETURN VALUES Upon successful completion, `posix_spawnattr_destroy()` and `posix_spawnattr_init()` return 0. Otherwise, an error number is returned to indicate the error.

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

- ENOMEM Insufficient memory exists to initialize the spawn attributes object.

The `posix_spawnattr_destroy()` function may fail if:

- EINVAL The value specified by `attr` is invalid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

posix_spawnattr_destroy(3C)

SEE ALSO | [posix_spawn\(3C\)](#), [posix_spawnattr_getflags\(3C\)](#),
[posix_spawnattr_getpgroup\(3C\)](#), [posix_spawnattr_getschedparam\(3C\)](#),
[posix_spawnattr_getschedpolicy\(3C\)](#),
[posix_spawnattr_getsigdefault\(3C\)](#), [posix_spawnattr_getsigmask\(3C\)](#),
[attributes\(5\)](#), [standards\(5\)](#)

NAME posix_spawnattr_getflags, posix_spawnattr_setflags – get and set spawn-flags attribute of spawn attributes object

SYNOPSIS #include <spawn.h>

```
int posix_spawnattr_getflags(const posix_spawnattr_t *restrict
    attr, short *restrict flags);
int posix_spawnattr_setflags(posix_spawnattr_t * attr, short flags);
```

DESCRIPTION The `posix_spawnattr_getflags()` function obtains the value of the *spawn-flags* attribute from the attributes object referenced by *attr*.

The `posix_spawnattr_setflags()` function sets the *spawn-flags* attribute in an initialized attributes object referenced by *attr*.

The *spawn-flags* attribute is used to indicate which process attributes are to be changed in the new process image when invoking `posix_spawn(3C)` or `posix_spawnp(3C)`. It is the bitwise-inclusive OR of zero or more of the following flags:

```
POSIX_SPAWN_RESETIDS
POSIX_SPAWN_SETPGROUP
POSIX_SPAWN_SETSIGDEF
POSIX_SPAWN_SETSIGMASK
POSIX_SPAWN_SETSCHEDPARAM
POSIX_SPAWN_SETSCHEDULER
```

These flags are defined in `<spawn.h>`. The default value of this attribute is as if no flags were set.

RETURN VALUES Upon successful completion, `posix_spawnattr_getflags()` returns 0 and stores the value of the *spawn-flags* attribute of *attr* into the object referenced by the *flags* parameter. Otherwise, an error number is returned to indicate the error.

Upon successful completion, `posix_spawnattr_setflags()` returns 0. Otherwise, an error number is returned to indicate the error.

ERRORS These functions may fail if:

EINVAL The value specified by *attr* is invalid.

The `posix_spawnattr_setflags()` function may fail if:

EINVAL The value of the attribute being set is not valid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

posix_spawnattr_getflags(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [posix_spawn\(3C\)](#), [posix_spawnattr_destroy\(3C\)](#),
[posix_spawnattr_getpgroup\(3C\)](#), [posix_spawnattr_getschedparam\(3C\)](#),
[posix_spawnattr_getschedpolicy\(3C\)](#),
[posix_spawnattr_getsigdefault\(3C\)](#), [posix_spawnattr_getsigmask\(3C\)](#),
[attributes\(5\)](#), [standards\(5\)](#)

NAME posix_spawnattr_getpgroup, posix_spawnattr_setpgroup – get and set spawn-pgroup attribute of spawn attributes object

SYNOPSIS #include <spawn.h>

```
int posix_spawnattr_getpgroup(const posix_spawnattr_t *restrict
    attr, pid_t *restrict pgroup);
int posix_spawnattr_setpgroup(posix_spawnattr_t *attr, pid_t
    pgroup);
```

DESCRIPTION The `posix_spawnattr_getpgroup()` function obtains the value of the `spawn-pgroup` attribute from the attributes object referenced by `attr`.

The `posix_spawnattr_setpgroup()` function sets the `spawn-pgroup` attribute in an initialized attributes object referenced by `attr`.

The `spawn-pgroup` attribute represents the process group to be joined by the new process image in a spawn operation (if `POSIX_SPAWN_SETPGROUP` is set in the `spawn-flags` attribute). The default value of this attribute is zero.

RETURN VALUES Upon successful completion, `posix_spawnattr_getpgroup()` returns 0 and stores the value of the `spawn-pgroup` attribute of `attr` into the object referenced by the `pgroup` parameter. Otherwise, an error number is returned to indicate the error.

Upon successful completion, `posix_spawnattr_setpgroup()` returns 0. Otherwise, an error number is returned to indicate the error.

ERRORS These functions may fail if:

EINVAL The value specified by `attr` is invalid.

The `posix_spawnattr_setpgroup()` function may fail if:

EINVAL The value of the attribute being set is not valid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `posix_spawn(3C)`, `posix_spawnattr_getpgroup(3C)`, `posix_spawnattr_getpgroup(3C)`, `posix_spawnattr_getschedparam(3C)`, `posix_spawnattr_getschedpolicy(3C)`, `posix_spawnattr_getsigdefault(3C)`, `posix_spawnattr_getsigmask(3C)`, `attributes(5)`, `standards(5)`

posix_spawnattr_getschedparam(3C)

NAME	posix_spawnattr_getschedparam, posix_spawnattr_setschedparam – get and set spawn-schedparam attribute of spawn attributes object						
SYNOPSIS	<pre>#include <spawn.h> #include <sched.h> int posix_spawnattr_getschedparam(const posix_spawnattr_t *restrict attr, struct sched_param *restrict schedparam); int posix_spawnattr_setschedparam(posix_spawnattr_t *restrict attr, const struct sched_param *restrict schedparam);</pre>						
DESCRIPTION	<p>The <code>posix_spawnattr_getschedparam()</code> function obtains the value of the <code>spawn-schedparam</code> attribute from the attributes object referenced by <code>attr</code>.</p> <p>The <code>posix_spawnattr_setschedparam()</code> function sets the <code>spawn-schedparam</code> attribute in an initialized attributes object referenced by <code>attr</code>.</p> <p>The <code>spawn-schedparam</code> attribute represents the scheduling parameters to be assigned to the new process image in a spawn operation (if <code>POSIX_SPAWN_SETSCHEDULER</code> or <code>POSIX_SPAWN_SETSCHEDPARAM</code> is set in the <code>spawn-flags</code> attribute). The default value of this attribute is unspecified.</p>						
RETURN VALUES	<p>Upon successful completion, <code>posix_spawnattr_getschedparam()</code> returns 0 and stores the value of the <code>spawn-schedparam</code> attribute of <code>attr</code> into the object referenced by the <code>schedparam</code> parameter. Otherwise, an error number is returned to indicate the error.</p> <p>Upon successful completion, <code>posix_spawnattr_setschedparam()</code> returns 0. Otherwise, an error number is returned to indicate the error.</p>						
ERRORS	<p>These functions may fail if:</p> <p>EINVAL The value specified by <code>attr</code> is invalid.</p> <p>The <code>posix_spawnattr_setschedparam()</code> function may fail if:</p> <p>EINVAL The value of the attribute being set is not valid.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>posix_spawn(3C)</code> , <code>posix_spawnattr_destroy(3C)</code> , <code>posix_spawnattr_getflags(3C)</code> , <code>posix_spawnattr_getpgroup(3C)</code> , <code>posix_spawnattr_getschedpolicy(3C)</code> , <code>posix_spawnattr_getsigdefault(3C)</code> , <code>posix_spawnattr_getsigmask(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME posix_spawnattr_getschedpolicy, posix_spawnattr_setschedpolicy – get and set spawn-schedpolicy attribute of spawn attributes object

SYNOPSIS

```
#include <spawn.h>
#include <sched.h>

int posix_spawnattr_getschedpolicy(const posix_spawnattr_t
    *restrict attr, int *restrict schedpolicy);

int posix_spawnattr_setschedpolicy(posix_spawnattr_t *attr, int
    schedpolicy);
```

DESCRIPTION

The `posix_spawnattr_getschedpolicy()` function obtains the value of the `spawn-schedpolicy` attribute from the attributes object referenced by `attr`.

The `posix_spawnattr_setschedpolicy()` function sets the `spawn-schedpolicy` attribute in an initialized attributes object referenced by `attr`.

The `spawn-schedpolicy` attribute represents the scheduling policy to be assigned to the new process image in a spawn operation (if `POSIX_SPAWN_SETSCHEDULER` is set in the `spawn-flags` attribute). The default value of this attribute is unspecified.

RETURN VALUES

Upon successful completion, `posix_spawnattr_getschedpolicy()` returns 0 and stores the value of the `spawn-schedpolicy` attribute of `attr` into the object referenced by the `schedpolicy` parameter. Otherwise, an error number is returned to indicate the error.

Upon successful completion, `posix_spawnattr_setschedpolicy()` returns 0. Otherwise, an error number is returned to indicate the error.

ERRORS

These functions may fail if:

`EINVAL` The value specified by `attr` is invalid.

The `posix_spawnattr_setschedpolicy()` function may fail if:

`EINVAL` The value of the attribute being set is not valid.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO

`posix_spawn(3C)`, `posix_spawnattr_destroy(3C)`, `posix_spawnattr_getflags(3C)`, `posix_spawnattr_getpgroup(3C)`, `posix_spawnattr_getschedparam(3C)`, `posix_spawnattr_getsigdefault(3C)`, `posix_spawnattr_getsigmask(3C)`, `attributes(5)`, `standards(5)`

posix_spawnattr_getsigdefault(3C)

NAME	posix_spawnattr_getsigdefault, posix_spawnattr_setsigdefault – get and set spawn-sigdefault attribute of spawn attributes object						
SYNOPSIS	<pre>#include <signal.h> #include <spawn.h> int posix_spawnattr_getsigdefault(const posix_spawnattr_t *restrict attr, sigset_t *restrict sigdefault); int posix_spawnattr_setsigdefault(posix_spawnattr_t *restrict attr, const sigset_t *restrict sigdefault);</pre>						
DESCRIPTION	<p>The <code>posix_spawnattr_getsigdefault()</code> function obtains the value of the <code>spawn-sigdefault</code> attribute from the attributes object referenced by <code>attr</code>.</p> <p>The <code>posix_spawnattr_setsigdefault()</code> function sets the <code>spawn-sigdefault</code> attribute in an initialized attributes object referenced by <code>attr</code>.</p> <p>The <code>spawn-sigdefault</code> attribute represents the set of signals to be forced to default signal handling in the new process image (if <code>POSIX_SPAWN_SETSIGDEF</code> is set in the <code>spawn-flags</code> attribute) by a spawn operation. The default value of this attribute is an empty signal set.</p>						
RETURN VALUES	<p>Upon successful completion, <code>posix_spawnattr_getsigdefault()</code> returns 0 and stores the value of the <code>spawn-sigdefault</code> attribute of <code>attr</code> into the object referenced by the <code>sigdefault</code> parameter. Otherwise, an error number is returned to indicate the error.</p> <p>Upon successful completion, <code>posix_spawnattr_setsigdefault()</code> returns 0. Otherwise, an error number is returned to indicate the error.</p>						
ERRORS	<p>These functions may fail if:</p> <p>EINVAL The value specified by <code>attr</code> is invalid.</p> <p>The <code>posix_spawnattr_setsigdefault()</code> function may fail if:</p> <p>EINVAL The value of the attribute being set is not valid.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>posix_spawn(3C)</code> , <code>posix_spawnattr_destroy(3C)</code> , <code>posix_spawnattr_getflags(3C)</code> , <code>posix_spawnattr_getpgroup(3C)</code> , <code>posix_spawnattr_getschedparam(3C)</code> , <code>posix_spawnattr_getschedpolicy(3C)</code> , <code>posix_spawnattr_getsigmask(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME posix_spawnattr_getsigmask, posix_spawnattr_setsigmask – get and set spawn-sigmask attribute of spawn attributes object

SYNOPSIS

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

int posix_spawnattr_getsigmask(const posix_spawnattr_t *restrict
    attr, sigset_t *restrict sigmask);

int posix_spawnattr_setsigmask(posix_spawnattr_t *restrict attr,
    const sigset_t *restrict sigmask);
```

DESCRIPTION

The `posix_spawnattr_getsigmask()` function obtains the value of the `spawn-sigmask` attribute from the attributes object referenced by `attr`.

The `posix_spawnattr_setsigmask()` function sets the `spawn-sigmask` attribute in an initialized attributes object referenced by `attr`.

The `spawn-sigmask` attribute represents the signal mask in effect in the new process image of a spawn operation (if `POSIX_SPAWN_SETSIGMASK` is set in the `spawn-flags` attribute). The default value of this attribute is unspecified.

RETURN VALUES

Upon successful completion, `posix_spawnattr_getsigmask()` returns 0 and stores the value of the `spawn-sigmask` attribute of `attr` into the object referenced by the `sigmask` parameter. Otherwise, an error number is returned to indicate the error.

Upon successful completion, `posix_spawnattr_setsigmask()` returns 0. Otherwise, an error number is returned to indicate the error.

ERRORS

These functions may fail if:

EINVAL The value specified by `attr` is invalid.

The `posix_spawnattr_setsigmask()` function may fail if:

EINVAL The value of the attribute being set is not valid.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `posix_spawn(3C)`, `posix_spawnattr_destroy(3C)`, `posix_spawnattr_getflags(3C)`, `posix_spawnattr_getpgroup(3C)`, `posix_spawnattr_getschedparam(3C)`, `posix_spawnattr_getschedpolicy(3C)`, `posix_spawnattr_getsigmask(3C)`, `attributes(5)`, `standards(5)`

posix_spawn_file_actions_addclose(3C)

NAME	posix_spawn_file_actions_addclose, posix_spawn_file_actions_addopen – add close or open action to spawn file actions object		
SYNOPSIS	<pre>#include <spawn.h> int posix_spawn_file_actions_addclose(posix_spawn_file_actions_t *file_actions, int <i>fildes</i>,); int posix_spawn_file_actions_addopen(posix_spawn_file_actions_t *restrict file_actions, int <i>fildes</i>, const char *restrict <i>path</i>, int <i>oflag</i>, mode_t <i>mode</i>);</pre>		
DESCRIPTION	<p>These functions add or delete a close or open action to a spawn file actions object.</p> <p>A spawn file actions object is of type <code>posix_spawn_file_actions_t</code> (defined in <code><spawn.h></code>) and is used to specify a series of actions to be performed by a <code>posix_spawn(3C)</code> or <code>posix_spawn(3C)</code> operation to arrive at the set of open file descriptors for the child process given the set of open file descriptors of the parent.</p> <p>A spawn file actions object, when passed to <code>posix_spawn()</code> or <code>posix_spawnp()</code>, specifies how the set of open file descriptors in the calling process is transformed into a set of potentially open file descriptors for the spawned process. This transformation occurs as though the specified sequence of actions was performed exactly once, in the context of the spawned process (prior to execution of the new process image), in the order in which the actions were added to the object. Additionally, when the new process image is executed, any file descriptor (from this new set) which has its <code>FD_CLOEXEC</code> flag set is closed (see <code>posix_spawn(3C)</code>).</p> <p>The <code>posix_spawn_file_actions_addclose()</code> function adds a close action to the object referenced by <code>file_actions</code> that causes the file descriptor <code>fildes</code> to be closed (as if <code>close(fildes)</code> had been called) when a new process is spawned using this file actions object.</p> <p>The <code>posix_spawn_file_actions_addopen()</code> function adds an open action to the object referenced by <code>file_actions</code> that causes the file named by <code>path</code> to be opened (as if <code>open(path, oflag, mode)</code> had been called, and the returned file descriptor, if not <code>fildes</code>, had been changed to <code>fildes</code>) when a new process is spawned using this file actions object. If <code>fildes</code> was already an open file descriptor, it is closed before the new file is opened.</p> <p>The string described by <code>path</code> is copied by the <code>posix_spawn_file_actions_addopen()</code> function.</p>		
RETURN VALUES	Upon successful completion, these functions return 0. Otherwise, an error number is returned to indicate the error.		
ERRORS	<p>These functions will fail if:</p> <table><tr><td>EBADF</td><td>The value specified by <code>fildes</code> is negative or greater than or equal to <code>{OPEN_MAX}</code>.</td></tr></table> <p>These functions may fail if:</p>	EBADF	The value specified by <code>fildes</code> is negative or greater than or equal to <code>{OPEN_MAX}</code> .
EBADF	The value specified by <code>fildes</code> is negative or greater than or equal to <code>{OPEN_MAX}</code> .		

`posix_spawn_file_actions_addclose(3C)`

`EINVAL` The value specified by *file_actions* is invalid.

`ENOMEM` Insufficient memory exists to add to the spawn file actions object.

It is not considered an error for the *fildevs* argument passed to these functions to specify a file descriptor for which the specified operation could not be performed at the time of the call. Any such error will be detected when the associated file actions object is later used during a `posix_spawn()` or `posix_spawnnp()` operation.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `close(2)`, `dup(2)`, `open(2)`, `posix_spawn(3C)`, `posix_spawn_file_actions_adddup2(3C)`, `posix_spawn_file_actions_destroy(3C)`, `attributes(5)`, `standards(5)`

posix_spawn_file_actions_adddup2(3C)

NAME	posix_spawn_file_actions_adddup2 – add dup2 action to spawn file actions object						
SYNOPSIS	<pre>#include <spawn.h> int posix_spawn_file_actions_adddup2(posix_spawn_file_actions_t *file_actions, int fildes, int newfildes);</pre>						
DESCRIPTION	<p>The <code>posix_spawn_file_actions_adddup2()</code> function adds a dup2(3C) action to the object referenced by <code>file_actions</code> that causes the file descriptor <code>fildes</code> to be duplicated as <code>newfildes</code> (as if <code>dup2(fildes, newfildes)</code> had been called) when a new process is spawned using this file actions object.</p> <p>A spawn file actions object is as defined in posix_spawn_file_actions_addclose(3C).</p>						
RETURN VALUES	Upon successful completion, the <code>posix_spawn_file_actions_adddup2()</code> function returns 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	<p>The <code>posix_spawn_file_actions_adddup2()</code> function will fail if:</p> <table><tr><td>EBADF</td><td>The value specified by <code>fildes</code> or <code>newfildes</code> is negative or greater than or equal to {OPEN_MAX}.</td></tr><tr><td>ENOMEM</td><td>Insufficient memory exists to add to the spawn file actions object.</td></tr></table> <p>The <code>posix_spawn_file_actions_adddup2()</code> function may fail if:</p> <table><tr><td>EINVAL</td><td>The value specified by <code>file_actions</code> is invalid.</td></tr></table> <p>It is not considered an error for the <code>fildes</code> argument passed to <code>posix_spawn_file_actions_adddup2()</code> to specify a file descriptor for which the specified operation could not be performed at the time of the call. Any such error will be detected when the associated file actions object is later used during a posix_spawn(3C) or posix_spawnnp(3C) operation.</p>	EBADF	The value specified by <code>fildes</code> or <code>newfildes</code> is negative or greater than or equal to {OPEN_MAX}.	ENOMEM	Insufficient memory exists to add to the spawn file actions object.	EINVAL	The value specified by <code>file_actions</code> is invalid.
EBADF	The value specified by <code>fildes</code> or <code>newfildes</code> is negative or greater than or equal to {OPEN_MAX}.						
ENOMEM	Insufficient memory exists to add to the spawn file actions object.						
EINVAL	The value specified by <code>file_actions</code> is invalid.						
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>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	dup2(3C) , posix_spawn(3C) , posix_spawn_file_actions_addclose(3C) , posix_spawn_file_actions_destroy(3C) , attributes(5) , standards(5)						

NAME posix_spawn_file_actions_destroy, posix_spawn_file_actions_init – destroy and initialize spawn file actions object

SYNOPSIS

```
#include <spawn.h>

int posix_spawn_file_actions_destroy(posix_spawn_file_actions_t
    *file_actions) ;

int posix_spawn_file_actions_init(posix_spawn_file_actions_t
    *file_actions) ;
```

DESCRIPTION The `posix_spawn_file_actions_destroy()` function destroys the object referenced by `file_actions`. The object becomes, in effect, uninitialized. An implementation can cause `posix_spawn_file_actions_destroy()` to set the object referenced by `file_actions` to an invalid value. A destroyed spawn file actions object can be reinitialized using `posix_spawn_file_actions_init()`. The results of otherwise referencing the object after it has been destroyed are undefined.

The `posix_spawn_file_actions_init()` function initializes the object referenced by `file_actions` to contain no file actions for `posix_spawn(3C)` or `posix_spawnnp(3C)` to perform.

A spawn file actions object is as defined in `posix_spawn_file_actions_addclose(3C)`.

The effect of initializing an already initialized spawn file actions object is undefined.

RETURN VALUES Upon successful completion, these functions return 0. Otherwise, an error number is returned to indicate the error.

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

ENOMEM Insufficient memory exists to initialize the spawn file actions object.

The `posix_spawn_file_actions_destroy()` function will may if:

EINVAL The value specified by `file_actions` is invalid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `posix_spawn(3C)`, `posix_spawn_file_actions_addclose(3C)`, `attributes(5)`, `standards(5)`

printf(3C)

NAME	printf, fprintf, sprintf, snprintf – print formatted output
SYNOPSIS	<pre>#include <stdio.h> int printf(const char *restrict format, /* args*/ ...); int fprintf(FILE *restrict stream, const char *restrict format, /* args*/ ...); int sprintf(char *restrict s, const char *restrict format, /* args*/ ...); int snprintf(char *restrict s, size_t n, const char *restrict format, /* args*/ ...);</pre>
DESCRIPTION	<p>The <code>printf()</code> function places output on the standard output stream <code>stdout</code>.</p> <p>The <code>fprintf()</code> function places output on on the named output stream <code>stream</code>.</p> <p>The <code>sprintf()</code> function places output, followed by the null byte (<code>\0</code>), in consecutive bytes starting at <code>s</code>; it is the user's responsibility to ensure that enough storage is available.</p> <p>The <code>snprintf()</code> function is identical to <code>sprintf()</code> with the addition of the argument <code>n</code>, which specifies the size of the buffer referred to by <code>s</code>. If <code>n</code> is 0, nothing is written and <code>s</code> can be a null pointer. Otherwise, output bytes beyond the <code>n</code>-1st are discarded instead of being written to the array and a null byte is written at the end of the bytes actually written into the array.</p> <p>Each of these functions converts, formats, and prints its arguments under control of the <i>format</i>. The <i>format</i> is a character string, beginning and ending in its initial shift state, if any. The <i>format</i> is composed of zero or more directives: <i>ordinary characters</i>, which are simply copied to the output stream and <i>conversion specifications</i>, each of which results in the fetching of zero or more arguments. The results are undefined if there are insufficient arguments for the <i>format</i>. If the <i>format</i> is exhausted while arguments remain, the excess arguments are evaluated but are otherwise ignored.</p> <p>Conversions can be applied to the <i>n</i>th argument after the <i>format</i> in the argument list, rather than to the next unused argument. In this case, the conversion specifier <code>%</code> (see below) is replaced by the sequence <code>%n\$</code>, where <code>n</code> is a decimal integer in the range <code>[1, NL_ARGMAX]</code>, giving the position of the argument in the argument list. This feature provides for the definition of format strings that select arguments in an order appropriate to specific languages (see the <code>EXAMPLES</code> section).</p> <p>In format strings containing the <code>%n\$</code> form of conversion specifications, numbered arguments in the argument list can be referenced from the format string as many times as required.</p> <p>In format strings containing the <code>%</code> form of conversion specifications, each argument in the argument list is used exactly once.</p>

Conversion Specifications

All forms of the `printf()` functions allow for the insertion of a language-dependent radix character in the output string. The radix character is defined by the program's locale (category `LC_NUMERIC`). In the POSIX locale, or in a locale where the radix character is not defined, the radix character defaults to a period (`.`).

Each conversion specification is introduced by the `%` character or by the character sequence `%n$`, after which the following appear in sequence:

- An optional field, consisting of a decimal digit string followed by a `$`, specifying the next argument to be converted. If this field is not provided, the *args* following the last argument converted will be used.
- Zero or more *flags* (in any order), which modify the meaning of the conversion specification.
- An optional minimum *field width*. If the converted value has fewer bytes than the field width, it will be padded with spaces by default on the left; it will be padded on the right, if the left-adjustment flag (`‐`), described below, is given to the field width. The field width takes the form of an asterisk (`*`), described below, or a decimal integer.

If the conversion specifier is *s*, a standard-conforming application (see `standards(5)`) interprets the field width as the minimum number of bytes to be printed; an application that is not standard-conforming interprets the field width as the minimum number of columns of screen display. For an application that is not standard-conforming, `%10s` means if the converted value has a screen width of 7 columns, 3 spaces would be padded on the right.

If the format is `%ws`, then the field width should be interpreted as the minimum number of columns of screen display.

- An optional *precision* that gives the minimum number of digits to appear for the *d*, *i*, *o*, *u*, *x*, and *X* conversions (the field is padded with leading zeros); the number of digits to appear after the radix character for the *a*, *A*, *e*, *E*, *f*, and *F* conversions, the maximum number of significant digits for the *g* and *G* conversions; or the maximum number of bytes to be printed from a string in *s* and *S* conversions. The precision takes the form of a period (`.`) followed either by an asterisk (`*`), described below, or an optional decimal digit string, where a null digit string is treated as 0. If a precision appears with any other conversion specifier, the behavior is undefined.

If the conversion specifier is *s* or *S*, a standard-conforming application (see `standards(5)`) interprets the precision as the maximum number of bytes to be written; an application that is not standard-conforming interprets the precision as the maximum number of columns of screen display. For an application that is not standard-conforming, `%.5s` would print only the portion of the string that would display in 5 screen columns. Only complete characters are written.

For `%ws`, the precision should be interpreted as the maximum number of columns of screen display. The precision takes the form of a period (`.`) followed by a decimal digit string; a null digit string is treated as zero. Padding specified by the precision overrides the padding specified by the field width.

- An optional *length modifier* that specified the size of the argument.
- A *conversion specifier* that indicates the type of conversion to be applied.

printf(3C)

A field width, or precision, or both can be indicated by an asterisk (*). In this case, an argument of type `int` supplies the field width or precision. Arguments specifying field width, or precision, or both must appear in that order before the argument, if any, to be converted. A negative field width is taken as a `-` flag followed by a positive field width. A negative precision is taken as if the precision were omitted. In format strings containing the `%n$` form of a conversion specification, a field width or precision may be indicated by the sequence `*m$`, where `m` is a decimal integer in the range [1, `NL_ARGMAX`] giving the position in the argument list (after the format argument) of an integer argument containing the field width or precision, for example:

```
printf("%1$d:%2$.*3$d:%4$.*3$d\n", hour, min, precision, sec);
```

The *format* can contain either numbered argument specifications (that is, `%n$` and `*m$`), or unnumbered argument specifications (that is, `%` and `*`), but normally not both. The only exception to this is that `%%` can be mixed with the `%n$` form. The results of mixing numbered and unnumbered argument specifications in a *format* string are undefined. When numbered argument specifications are used, specifying the *N*th argument requires that all the leading arguments, from the first to the (*N*-1)th, are specified in the format string.

Flag Characters

The flag characters and their meanings are:

- ' The integer portion of the result of a decimal conversion (`%i`, `%d`, `%u`, `%f`, `%F`, `%g`, or `%G`) will be formatted with thousands' grouping characters. For other conversions the behavior is undefined. The non-monetary grouping character is used.
- The result of the conversion will be left-justified within the field. The conversion will be right-justified if this flag is not specified.
- + The result of a signed conversion will always begin with a sign (+ or -). The conversion will begin with a sign only when a negative value is converted if this flag is not specified.
- space If the first character of a signed conversion is not a sign or if a signed conversion results in no characters, a space will be placed before the result. This means that if the `space` and `+` flags both appear, the `space` flag will be ignored.
- # The value is to be converted to an alternate form. For `c`, `d`, `i`, `s`, and `u` conversions, the flag has no effect. For an `o` conversion, it increases the precision (if necessary) to force the first digit of the result to be a zero. For `x` or `X` conversion, a non-zero result will have `0x` (or `0X`) prepended to it. For `a`, `A`, `e`, `E`, `f`, `F`, `g`, and `G` conversions, the result will always contain a radix character, even if no digits follow the radix character. Without this flag, the radix character appears in the result of these conversions only if a digit follows it. For `g` and `G` conversions, trailing zeros will not be removed from the result as they normally are.
- 0 For `d`, `i`, `o`, `u`, `x`, `X`, `a`, `A`, `e`, `E`, `f`, `F`, `g`, and `G` conversions, leading zeros (following any indication of sign or base) are used to pad to the field width; no space padding is performed. If the `0` and `-` flags both appear, the

0 flag will be ignored. For `d`, `i`, `o`, `u`, `x`, and `X` conversions, if a precision is specified, the 0 flag will be ignored. If the 0 and ' flags both appear, the grouping characters are inserted before zero padding. For other conversions, the behavior is undefined.

Length Modifiers

The length modifiers and their meanings are:

hh	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , or <code>X</code> conversion specifier applies to a <code>signed char</code> or <code>unsigned char</code> argument (the argument will have been promoted according to the integer promotions, but its value will be converted to <code>signed char</code> or <code>unsigned char</code> before printing); or that a following <code>n</code> conversion specifier applies to a pointer to a <code>signed char</code> argument.
h	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , or <code>X</code> conversion specifier applies to a <code>short</code> or <code>unsigned short</code> argument (the argument will have been promoted according to the integer promotions, but its value will be converted to <code>short</code> or <code>unsigned short</code> before printing); or that a following <code>n</code> conversion specifier applies to a pointer to a <code>short</code> argument.
l (ell)	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , or <code>X</code> conversion specifier applies to a <code>long</code> or <code>unsigned long</code> argument; that a following <code>n</code> conversion specifier applies to a pointer to a <code>long</code> argument; that a following <code>c</code> conversion specifier applies to a <code>wint_t</code> argument; that a following <code>s</code> conversion specifier applies to a pointer to a <code>wchar_t</code> argument; or has no effect on a following <code>a</code> , <code>A</code> , <code>e</code> , <code>E</code> , <code>f</code> , <code>F</code> , <code>g</code> , or <code>G</code> conversion specifier.
ll (ell-ell)	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , or <code>X</code> conversion specifier applies to a <code>long long</code> or <code>unsigned long long</code> argument; or that a following <code>n</code> conversion specifier applies to a pointer to a <code>long long</code> argument.
j	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , or <code>X</code> conversion specifier applies to an <code>intmax_t</code> or <code>uintmax_t</code> argument; or that a following <code>n</code> conversion specifier applies to a pointer to an <code>intmax_t</code> argument. See NOTES.
z	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , or <code>X</code> conversion specifier applies to a <code>size_t</code> or the corresponding signed integer type argument; or that a following <code>n</code> conversion specifier applies to a pointer to a signed integer type corresponding to <code>size_t</code> argument.
t	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , or <code>X</code> conversion specifier applies to a <code>ptrdiff_t</code> or the corresponding unsigned type argument; or that a following <code>n</code> conversion specifier applies to a pointer to a <code>ptrdiff_t</code> argument.

printf(3C)

Conversion Specifiers

L	Specifies that a following <code>a</code> , <code>A</code> , <code>e</code> , <code>E</code> , <code>f</code> , <code>F</code> , <code>g</code> , or <code>G</code> conversion specifier applies to a <code>long double</code> argument.
	If a length modifier appears with any conversion specifier other than as specified above, the behavior is undefined.
	Each conversion specifier results in fetching zero or more arguments. The results are undefined if there are insufficient arguments for the format. If the format is exhausted while arguments remain, the excess arguments are ignored.
	The conversion specifiers and their meanings are:
<code>d</code> , <code>i</code>	The <code>int</code> argument is converted to a signed decimal in the style <code>[-]ddd</code> . The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no characters.
<code>o</code>	The unsigned <code>int</code> argument is converted to unsigned octal format in the style <code>ddd</code> . The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no characters.
<code>u</code>	The unsigned <code>int</code> argument is converted to unsigned decimal format in the style <code>ddd</code> . The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no characters.
<code>x</code>	The unsigned <code>int</code> argument is converted to unsigned hexadecimal format in the style <code>ddd</code> ; the letters <code>abcdef</code> are used. The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeros. The default precision is 1. The result of converting 0 with an explicit precision of 0 is no characters.
<code>X</code>	Behaves the same as the <code>x</code> conversion specifier except that letters <code>ABCDEF</code> are used instead of <code>abcdef</code> .
<code>f</code> , <code>F</code>	The <code>double</code> argument is converted to decimal notation in the style <code>[-]ddd.ddd</code> , where the number of digits after the radix character (see setlocale(3C)) is equal to the precision specification. If the precision is missing it is taken as 6; if the precision is explicitly 0 and the <code>#</code> flag is not specified, no radix character appears. If a radix character appears, at least 1 digit appears before it. The converted value is rounded to fit the specified output format according to the prevailing floating point rounding direction mode. If the conversion is not exact, an inexact exception is raised.

For the `f` specifier, a double argument representing an infinity or NaN is converted in the style of the `e` conversion specifier, except that for an infinite argument, "infinity" or "Infinity" is printed when the precision is at least 8 and "inf" or "Inf" is printed otherwise.

For the `F` specifier, a double argument representing an infinity or NaN is converted in the SUSv3 style of the `E` conversion specifier, except that for an infinite argument, "INFINITY" is printed when the precision is at least 8 and or "INF" is printed otherwise.

`e, E` The double argument is converted to the style `[-]d.ddde±dd`, where there is one digit before the radix character (which is non-zero if the argument is non-zero) and the number of digits after it is equal to the precision. When the precision is missing it is taken as 6; if the precision is 0 and the `#` flag is not specified, no radix character appears. The `E` conversion specifier will produce a number with `E` instead of `e` introducing the exponent. The exponent always contains at least two digits. The converted value is rounded to fit the specified output format according to the prevailing floating point rounding direction mode. If the conversion is not exact, an inexact exception is raised.

Infinity and NaN values are handled in one of the following ways:

`SUSv3` For the `e` specifier, a double argument representing an infinity is printed as "`[-]infinity`", when the precision for the conversion is at least 7 and as "`[-]inf`" otherwise. A double argument representing a NaN is printed as "`[-]nan`". For the `E` specifier, "INF", "INFINITY", and "NaN" are printed instead of "inf", "infinity", and "nan", respectively. Printing of the sign follows the rules described above.

`Default` A double argument representing an infinity is printed as "`[-]Infinity`", when the precision for the conversion is at least 7 and as "`[-]Inf`" otherwise. A double argument representing a NaN is printed as "`[-]NaN`". Printing of the sign follows the rules described above.

`g, G` The double argument is printed in style `f` or `e` (or in style `E` in the case of a `G` conversion specifier), with the precision specifying the number of significant digits. If an explicit precision is 0, it is taken as 1. The style used depends on the value converted: style `e` (or `E`) will be used only if the exponent resulting from the conversion is less than `-4` or greater than or equal to the precision. Trailing zeros are removed from the fractional part of the result. A radix character appears only if it is followed by a digit.

printf(3C)

	<p>A double argument representing an infinity or NaN is converted in the style of the e or E conversion specifier, except that for an infinite argument, "infinity", "INFINITY", or "Infinity" is printed when the precision is at least 8 and "inf", "INF", or "Inf" is printed otherwise.</p>
a, A	<p>A double argument representing a floating-point number is converted in the style "[<i>-]0xh.hhhhp±d</i>", where the single hexadecimal digit preceding the radix point is 0 if the value converted is zero and 1 otherwise and the number of hexadecimal digits after it is equal to the precision; if the precision is missing, the number of digits printed after the radix point is 13 for the conversion of a double value, 16 for the conversion of a long double value on x86, and 28 for the conversion of a long double value on SPARC; if the precision is zero and the '#' flag is not specified, no decimal-point character will appear. The letters "abcdef" are used for a conversion and the letters "ABCDEF" for A conversion. The A conversion specifier produces a number with 'X' and 'P' instead of 'x' and 'p'. The exponent will always contain at least one digit, and only as many more digits as necessary to represent the decimal exponent of 2. If the value is zero, the exponent is zero.</p> <p>The converted value is rounded to fit the specified output format according to the prevailing floating point rounding direction mode. If the conversion is not exact, an inexact exception is raised.</p> <p>A double argument representing an infinity or NaN is converted in the SUSv3 style of an e or E conversion specifier.</p>
c	<p>The int argument is converted to an unsigned char, and the resulting byte is printed.</p> <p>If an l (ell) qualifier is present, the wint_t argument is converted as if by an ls conversion specification with no precision and an argument that points to a two-element array of type wchar_t, the first element of which contains the wint_t argument to the ls conversion specification and the second element contains a null wide-character.</p>
C	<p>Same as lc.</p>
wc	<p>The int argument is converted to a wide character (wchar_t), and the resulting wide character is printed.</p>
s	<p>The argument must be a pointer to an array of char. Bytes from the array are written up to (but not including) any terminating null byte. If a precision is specified, a standard-conforming application (see standards(5)) will write only the number of bytes specified by precision; an application that is not standard-conforming will write only the portion of the string that will display in the number of columns of screen display specified by precision. If the precision is not specified, it is taken to be infinite, so all bytes up to the first null byte are printed. An argument with a null value will yield undefined results.</p>

If an `l` (`ell`) qualifier is present, the argument must be a pointer to an array of type `wchar_t`. Wide-characters from the array are converted to characters (each as if by a call to the `wcrtomb(3C)` function, with the conversion state described by an `mbstate_t` object initialized to zero before the first wide-character is converted) up to and including a terminating null wide-character. The resulting characters are written up to (but not including) the terminating null character (byte). If no precision is specified, the array must contain a null wide-character. If a precision is specified, no more than that many characters (bytes) are written (including shift sequences, if any), and the array must contain a null wide-character if, to equal the character sequence length given by the precision, the function would need to access a wide-character one past the end of the array. In no case is a partial character written.

<code>s</code>	Same as <code>l s</code> .
<code>ws</code>	The argument must be a pointer to an array of <code>wchar_t</code> . Bytes from the array are written up to (but not including) any terminating null character. If the precision is specified, only that portion of the wide-character array that will display in the number of columns of screen display specified by precision will be written. If the precision is not specified, it is taken to be infinite, so all wide characters up to the first null character are printed. An argument with a null value will yield undefined results.
<code>p</code>	The argument must be a pointer to <code>void</code> . The value of the pointer is converted to a set of sequences of printable characters, which should be the same as the set of sequences that are matched by the <code>%p</code> conversion of the <code>scanf(3C)</code> function.
<code>n</code>	The argument must be a pointer to an integer into which is written the number of bytes written to the output standard I/O stream so far by this call to one of the <code>printf()</code> functions. No argument is converted.
<code>%</code>	Print a <code>%</code> ; no argument is converted. The entire conversion specification must be <code>%%</code> .

If a conversion specification does not match one of the above forms, the behavior is undefined.

In no case does a non-existent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is simply expanded to contain the conversion result. Characters generated by `printf()` and `fprintf()` are printed as if the `putc(3C)` function had been called.

The `st_ctime` and `st_mtime` fields of the file will be marked for update between the call to a successful execution of `printf()` or `fprintf()` and the next successful completion of a call to `fflush(3C)` or `fclose(3C)` on the same stream or a call to `exit(3C)` or `abort(3C)`.

printf(3C)

RETURN VALUES

The `printf()`, `fprintf()`, and `sprintf()` functions return the number of bytes transmitted (excluding the terminating null byte in the case of `sprintf()`).

The `snprintf()` function returns the number of bytes that would be written to `s` had `n` been sufficiently large excluding the terminating null byte. If the value of `n` is 0 on a call to `snprintf()`, the number of bytes that would have been written had `n` been sufficiently large excluding the terminating null are returned and `s` can be a null pointer.

Each function returns a negative value if an output error was encountered.

ERRORS

For the conditions under which `printf()` and `fprintf()` will fail and may fail, refer to [fputc\(3C\)](#) or [fputwc\(3C\)](#).

In addition, all forms of `printf()` may fail if:

`EILSEQ` A wide-character code that does not correspond to a valid character has been detected.

`EINVAL` There are insufficient arguments.

In addition, `printf()` and `fprintf()` may fail if:

`ENOMEM` Insufficient storage space is available.

USAGE

If the application calling the `printf()` functions has any objects of type `wint_t` or `wchar_t`, it must also include the header `<wchar.h>` to have these objects defined.

Escape Character Sequences

It is common to use the following escape sequences built into the C language when entering format strings for the `printf()` functions, but these sequences are processed by the C compiler, not by the `printf()` function.

`\a` Alert. Ring the bell.

`\b` Backspace. Move the printing position to one character before the current position, unless the current position is the start of a line.

`\f` Form feed. Move the printing position to the initial printing position of the next logical page.

`\n` Newline. Move the printing position to the start of the next line.

`\r` Carriage return. Move the printing position to the start of the current line.

`\t` Horizontal tab. Move the printing position to the next implementation-defined horizontal tab position on the current line.

`\v` Vertical tab. Move the printing position to the start of the next implementation-defined vertical tab position.

In addition, the C language supports character sequences of the form

`\octal-numberand`

\hex-number which translates into the character represented by the octal or hexadecimal number. For example, if ASCII representations are being used, the letter 'a' may be written as '\141' and 'Z' as '\132'. This syntax is most frequently used to represent the null character as '\0'. This is exactly equivalent to the numeric constant zero (0). Note that the octal number does not include the zero prefix as it would for a normal octal constant. To specify a hexadecimal number, omit the zero so that the prefix is an 'x' (uppercase 'X' is not allowed in this context). Support for hexadecimal sequences is an ANSI extension. See standards(5).

EXAMPLES

EXAMPLE 1 To print the language-independent date and time format, the following statement could be used:

```
printf (format, weekday, month, day, hour, min);
```

For American usage, *format* could be a pointer to the string:

```
"%s, %s %d, %d:%.2d\n"
```

producing the message:

```
Sunday, July 3, 10:02
```

whereas for German usage, *format* could be a pointer to the string:

```
"%1$s, %3$d. %2$s, %4$d:%5$.2d\n"
```

producing the message:

```
Sonntag, 3. Juli, 10:02
```

EXAMPLE 2 To print a date and time in the form *Sunday, July 3, 10:02*, where *weekday* and *month* are pointers to null-terminated strings:

```
printf("%s, %s %i, %d:%.2d", weekday, month, day, hour, min);
```

EXAMPLE 3 To print pi to 5 decimal places:

```
printf("pi = %.5f", 4 * atan(1.0));
```

Default

EXAMPLE 4 The following example applies only to applications which are not standard-conforming (see standards(5)). To print a list of names in columns which are 20 characters wide:

```
printf("%20s%20s%20s", lastname, firstname, middlename);
```

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	See below.

printf(3C)

The `sprintf()` and `snprintf()` functions are Async-Signal-Safe. The `printf()` and `fprintf()` functions can be used safely in multithreaded applications, as long as `setlocale(3C)` is not being called to change the locale.

SEE ALSO `exit(2)`, `lseek(2)`, `write(2)`, `abort(3C)`, `ecvt(3C)`, `exit(3C)`, `fclose(3C)`, `fflush(3C)`, `fputwc(3C)`, `putc(3C)`, `scanf(3C)`, `setlocale(3C)`, `stdio(3C)`, `vprintf(3C)`, `wcstombs(3C)`, `wctomb(3C)`, `attributes(5)`, `environ(5)`, `standards(5)`

NOTES If the `j` length modifier is used, 32-bit applications that were compiled using `c89` on releases prior to Solaris 10 will experience undefined behavior.

The `snprintf()` return value when $n = 0$ was changed in the Solaris 10 release. The change was based on the SUSv3 specification. The previous behavior was based on the initial SUSv2 specification, where `snprintf()` when $n = 0$ returns an unspecified value less than 1.

NAME	printf, fprintf, sprintf, vprintf, fprintf, vsprintf – formatted output conversion
SYNOPSIS	<pre> /usr/ucb/cc [flag ...] file ... #include <stdio.h> int printf(format, ...); const char *format; int fprintf(stream, format, va_list); FILE *stream; char *format; va_dcl; char *sprintf(s, format, va_list); char *s, *format; va_dcl; int vprintf(format, ap); char *format; va_list ap; int fprintf(stream, format, ap); FILE *stream; char *format; va_list ap; char *vsprintf(s, format, ap); char *s, *format; va_list ap; </pre>
DESCRIPTION	<p>printf() places output on the standard output stream stdout. fprintf() places output on the named output stream. sprintf() places “output,” followed by the NULL character (\0), in consecutive bytes starting at *s; it is the user’s responsibility to ensure that enough storage is available.</p> <p>vprintf(), fprintf(), and vsprintf() are the same as printf(), fprintf(), and sprintf() respectively, except that instead of being called with a variable number of arguments, they are called with an argument list as defined by <varargs.h>.</p> <p>Each of these functions converts, formats, and prints its args under control of the format. The format is a character string which contains two types of objects: plain characters, which are simply copied to the output stream, and conversion specifications, each of which causes conversion and printing of zero or more args. The results are undefined if there are insufficient args for the format. If the format is exhausted while args remain, the excess args are simply ignored.</p> <p>Each conversion specification is introduced by the character %. After the %, the following appear in sequence:</p> <ul style="list-style-type: none"> ■ Zero or more flags, which modify the meaning of the conversion specification.

printf(3UCB)

- An optional decimal digit string specifying a minimum *field width*. If the converted value has fewer characters than the field width, it will be padded on the left (or right, if the left-adjustment flag '-', described below, has been given) to the field width. The padding is with blanks unless the field width digit string starts with a zero, in which case the padding is with zeros.
- A *precision* that gives the minimum number of digits to appear for the d, i, o, u, x, or X conversions, the number of digits to appear after the decimal point for the e, E, and f conversions, the maximum number of significant digits for the g and G conversion, or the maximum number of characters to be printed from a string in s conversion. The precision takes the form of a period (.) followed by a decimal digit string; a NULL digit string is treated as zero. Padding specified by the precision overrides the padding specified by the field width.
- An optional l (ell) specifying that a following d, i, o, u, x, or X conversion character applies to a long integer *arg*. An l before any other conversion character is ignored.
- A character that indicates the type of conversion to be applied.

A field width or precision or both may be indicated by an asterisk (*) instead of a digit string. In this case, an integer *arg* supplies the field width or precision. The *arg* that is actually converted is not fetched until the conversion letter is seen, so the *args* specifying field width or precision must appear *before* the *arg* (if any) to be converted. A negative field width argument is taken as a '-' flag followed by a positive field width. If the precision argument is negative, it will be changed to zero.

The flag characters and their meanings are:

-	The result of the conversion will be left-justified within the field.
+	The result of a signed conversion will always begin with a sign (+ or -).
blank	If the first character of a signed conversion is not a sign, a blank will be prefixed to the result. This implies that if the blank and + flags both appear, the blank flag will be ignored.
#	This flag specifies that the value is to be converted to an "alternate form." For c, d, i, s, and u conversions, the flag has no effect. For o conversion, it increases the precision to force the first digit of the result to be a zero. For x or X conversion, a non-zero result will have 0x or 0X prefixed to it. For e, E, f, g, and G conversions, the result will always contain a decimal point, even if no digits follow the point (normally, a decimal point appears in the result of these conversions only if a digit follows it). For g and G conversions, trailing zeroes will <i>not</i> be removed from the result (which they normally are).

The conversion characters and their meanings are:

d,i,o,u,x,X	The integer <i>arg</i> is converted to signed decimal (d or i), unsigned octal (o), unsigned decimal (u), or unsigned hexadecimal notation
-------------	--

	(x and X), respectively; the letters abcdef are used for x conversion and the letters ABCDEF for X conversion. The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeroes. (For compatibility with older versions, padding with leading zeroes may alternatively be specified by prepending a zero to the field width. This does not imply an octal value for the field width.) The default precision is 1. The result of converting a zero value with a precision of zero is a NULL string.
f	The float or double <i>arg</i> is converted to decimal notation in the style <code>[-]ddd.ddd</code> where the number of digits after the decimal point is equal to the precision specification. If the precision is missing, 6 digits are given; if the precision is explicitly 0, no digits and no decimal point are printed.
e,E	The float or double <i>arg</i> is converted in the style <code>[-]d.ddd[±]ddd</code> , where there is one digit before the decimal point and the number of digits after it is equal to the precision; when the precision is missing, 6 digits are produced; if the precision is zero, no decimal point appears. The E format code will produce a number with E instead of e introducing the exponent. The exponent always contains at least two digits.
g,G	The float or double <i>arg</i> is printed in style f or e (or in style E in the case of a G format code), with the precision specifying the number of significant digits. The style used depends on the value converted: style e or E will be used only if the exponent resulting from the conversion is less than -4 or greater than the precision. Trailing zeroes are removed from the result; a decimal point appears only if it is followed by a digit.
	The e, E f, g, and G formats print IEEE indeterminate values (infinity or not-a-number) as "Infinity" or "NaN" respectively.
c	The character <i>arg</i> is printed.
s	The <i>arg</i> is taken to be a string (character pointer) and characters from the string are printed until a NULL character (<code>\0</code>) is encountered or until the number of characters indicated by the precision specification is reached. If the precision is missing, it is taken to be infinite, so all characters up to the first NULL character are printed. A NULL value for <i>arg</i> will yield undefined results.
%	Print a %; no argument is converted.

printf(3UCB)

In no case does a non-existent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is simply expanded to contain the conversion result. Padding takes place only if the specified field width exceeds the actual width. Characters generated by `printf()` and `fprintf()` are printed as if `putc(3C)` had been called.

RETURN VALUES Upon success, `printf()` and `fprintf()` return the number of characters transmitted, excluding the null character. `vprintf()` and `vfprintf()` return the number of characters transmitted. `sprintf()` and `vsprintf()` always return `s`. If an output error is encountered, `printf()`, `fprint()`, `vprintf()`, and `vfprintf()` return EOF.

EXAMPLES **EXAMPLE 1** Examples of the `printf` Command To Print a Date and Time

To print a date and time in the form "Sunday, July 3, 10:02," where *weekday* and *month* are pointers to NULL-terminated strings:

```
printf("%s, %s %i, %d:%.2d", weekday, month, day, hour, min);
```

EXAMPLE 2 Examples of the `printf` Command To Print to Five Decimal Places

To print to five decimal places:

```
printf("pi = %.5f", 4 * atan(1. 0));
```

SEE ALSO [econvert\(3C\)](#), [putc\(3C\)](#), [scanf\(3C\)](#), [vprintf\(3C\)](#)

NOTES Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.

Very wide fields (>128 characters) fail.

NAME	priv_addset, priv_allocset, priv_copyset, priv_delset, priv_emptyset, priv_fillset, priv_freeset, priv_intersect, priv_inverse, priv_isemptyset, priv_isequalset, priv_isfullset, priv_ismember, priv_issubset, priv_union – privilege set manipulation functions
SYNOPSIS	<pre>#include <priv.h> int priv_addset(priv_set_t *sp, const char *priv); priv_set_t *priv_allocset(void); void priv_copyset(const priv_set_t *src, priv_set_t *dst); int priv_delset(priv_set_t *sp, const char *priv); void priv_emptyset(priv_set_t *sp); void priv_fillset(priv_set_t *sp); void priv_freeset(priv_set_t *sp); void priv_intersect(const priv_set_t *src, priv_set_t *dst); void priv_inverse(priv_set_t *sp); boolean_t priv_isemptyset(const priv_set_t *sp); boolean_t priv_isequalset(const priv_set_t *src, const priv_set_t *dst); boolean_t priv_isfullset(const priv_set_t *sp); boolean_t priv_ismember(const priv_set_t *sp, const char *priv); boolean_t priv_issubset(const priv_set_t *src, const priv_set_t *dst); void priv_union(const priv_set_t *src, priv_set_t *dst);</pre>
DESCRIPTION	<p>The <i>sp</i>, <i>src</i>, and <i>dst</i> arguments point to privilege sets. The <i>priv</i> argument points to a named privilege.</p> <p>The <code>priv_addset()</code> function adds the named privilege <i>priv</i> to <i>sp</i>.</p> <p>The <code>priv_allocset()</code> function allocates sufficient memory to contain a privilege set. The value of the returned privilege set is indeterminate. The function returns NULL and sets <code>errno</code> when it fails to allocate memory.</p> <p>The <code>priv_copyset()</code> function copies the set <i>src</i> to <i>dst</i>.</p> <p>The <code>priv_delset()</code> function removes the named privilege <i>priv</i> from <i>sp</i>.</p> <p>The <code>priv_emptyset()</code> function clears all privileges from <i>sp</i>.</p> <p>The <code>priv_fillset()</code> function asserts all privileges in <i>sp</i>, including the privileges not currently defined in the system.</p>

priv_addset(3C)

The `priv_freeset()` function frees the storage allocated by `priv_allocset()`.

The `priv_intersect()` function intersects `src` with `dst` and places the results in `dst`.

The `priv_inverse()` function inverts the privilege set given as argument in place.

The `priv_isemptyset()` function checks whether the argument is an empty set.

The `priv_isequalset()` function checks whether the privilege set `src` is equal to `dst`.

The `priv_isfullset()` function checks whether the argument is a full set. A full set is a set with all bits set, regardless of whether the privilege is currently defined in the system.

The `priv_ismember()` function checks whether the named privilege `priv` is a member of `sp`.

The `priv_issubset()` function checks whether `src` is a subset of `dst`.

The `priv_union()` function takes the union of `src` and `dst` and places the result in `dst`.

RETURN VALUES

Upon successful completion, `priv_allocset()` returns a pointer to an opaque data structure. It returns `NULL` if memory allocation fails and sets `errno` to indicate the error.

Upon successful completion, `priv_isemptyset()`, `priv_isfullset()`, `priv_isequalset()`, `priv_issubset()`, and `priv_ismember()` return `B_TRUE`. Otherwise, they return `B_FALSE`.

Upon successful completion, `priv_delset()` and `priv_addset()` return 0. Otherwise, they return -1 and set `errno` to indicate the error.

ERRORS

The `priv_allocset()` function will fail if:

`ENOMEM` The physical limits of the system are exceeded by the memory allocation needed to hold a privilege set.

`EAGAIN` There is insufficient memory for allocation to hold a privilege set. The application can try again later.

The `priv_delset()` and `priv_addset()` functions will fail if:

`EINVAL` The privilege argument is not a valid privilege name.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

priv_addset(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [malloc\(3C\)](#), [attributes\(5\)](#), [privileges\(5\)](#)

NOTES The functions that compare sets operate on all bits of the set, regardless of whether the specific privileges are currently defined in the system.

priv_set(3C)

NAME priv_set, priv_ineffect – change privilege sets and check whether privileges are set

SYNOPSIS

```
#include <priv.h>

int priv_set(priv_op_t op, priv_ptype_t which, ...);

boolean_t priv_ineffect(const char *priv);
```

DESCRIPTION

The `priv_set()` function is a convenient wrapper for the `setppriv(2)` function. It takes three or more arguments. The operation argument, *op*, can be one of `PRIV_OFF`, `PRIV_ON` or `PRIV_SET`. The *which* argument is the name of the privilege set to change. The third argument is a list of zero or more privilege names terminated with a null pointer. If *which* is the special pseudo set `PRIV_ALLSETS`, the operation should be applied to all privilege sets.

The specified privileges are converted to a binary privilege set and `setppriv()` is called with the same *op* and *which* arguments. When called with `PRIV_ALLSETS` as the value for the *which* argument, `setppriv()` is called for each set in turn, aborting on the first failed call.

The `priv_ineffect()` function is a convenient wrapper for the `getppriv(2)` function. The *priv* argument specifies the name of the privilege for which this function checks its presence in the effective set.

RETURN VALUES

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

If *priv* is a valid privilege that is a member of the effective set, `priv_ineffect()` returns `B_TRUE`. Otherwise, it returns `B_FALSE` and sets `errno` to indicate the error.

ERRORS

The `priv_set()` function will fail if:

<code>EINVAL</code>	The value of <i>op</i> or <i>which</i> is out of range.
<code>ENOMEM</code>	Insufficient memory was allocated.
<code>EPERM</code>	The application attempted to add privileges to <code>PRIV_LIMIT</code> or <code>PRIV_PERMITTED</code> , or the application attempted to add privileges to <code>PRIV_INHERITABLE</code> or <code>PRIV_EFFECTIVE</code> that were not in <code>PRIV_PERMITTED</code> .

The `priv_ineffect()` function will fail if:

<code>EINVAL</code>	The privilege specified by <i>priv</i> is invalid.
<code>ENOMEM</code>	Insufficient memory was allocated.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving

priv_set(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [setppriv\(2\)](#), [priv_str_to_set\(3C\)](#), [attributes\(5\)](#), [privileges\(5\)](#)

priv_str_to_set(3C)

NAME	priv_str_to_set, priv_set_to_str, priv_getbyname, priv_getbynum, priv_getsetbyname, priv_getsetbynum, priv_gettext – privilege name functions
SYNOPSIS	<pre>#include <priv.h> priv_set_t *priv_str_to_set(const char *buf, const char *sep, const char **endptr); char *priv_set_to_str(const priv_set_t *set, char sep, int flag); int priv_getbyname(const char *privname); const char *priv_getbynum(int privnum); int priv_getsetbyname(const char *privsetname); const char *priv_getsetbynum(int privname); char *priv_gettext(const char *privname);</pre>
DESCRIPTION	<p>The <code>priv_str_to_set()</code> function maps the privilege specification in <i>buf</i> to a privilege set. It returns a privilege set on success or NULL on failure. If an error occurs when parsing the string, a pointer to the remainder of the string is stored in the object pointed to by <i>endptr</i>, provided that <i>endptr</i> is not a null pointer. If an error occurs when allocating memory, <code>errno</code> is set and the object pointed to by <i>endptr</i> is set to the null pointer, provided that <i>endptr</i> is not a null pointer.</p> <p>The application is responsible for freeing the returned privilege set using priv_freeset(3C).</p> <p>A privilege specification should contain one or more privilege names, separated by characters in <i>sep</i> using the same algorithm as strtok(3C). Privileges can optionally be preceded by a dash (-) or an exclamation mark (!), in which case they are excluded from the resulting set. The special strings “none” for the empty set, “all” for the set of all privileges, “zone” for the set of all privileges available within the caller’s zone, and “basic” for the set of basic privileges are also recognized. Set specifications are interpreted from left to right.</p> <p>The <code>priv_set_to_str()</code> function converts the privilege set set to a sequence of privileges separated by <i>sep</i>, returning the a pointer to the dynamically allocated result. The application is responsible for freeing the memory using free(3C).</p> <p>To maintain future compatibility, the “basic” set of privileges is included as “basic,!missing_basic_priv1,...”. When further currently unprivileged operations migrate to the basic privilege set, the conversion back of the result with <code>priv_str_to_set()</code> includes the additional basic privileges, guaranteeing that the resulting privilege set carries the same privileges. This behavior is the default and is equivalent to specifying a <i>flag</i> argument of <code>PRIV_STR_PORT</code>. When specifying a <i>flag</i> argument of <code>PRIV_STR_LIT</code>, the result does not treat basic privileges differently and the privileges present are all literally presented in the output. A flag argument of <code>PRIV_STR_SHORT</code> attempts to arrive at the shortest output, using the tokens “basic”, “zone”, “all”, and negated privileges. This output is most useful for trace output.</p>

The `priv_getbyname()` and `priv_getsetbyname()` functions map privilege names and privilege set names to numbers. The numbers returned are valid for the current kernel instance only and could change at the next boot. Only the privilege names should be committed to persistent storage. The numbers should not be committed to persistent storage. Both functions return -1 on error, setting `errno` to `EINVAL`.

The `priv_getbynum()` and `priv_getsetbynum()` functions map privileges numbers to names. The strings returned point to shared storage that should not be modified and is valid for the lifetime of the process. Both functions return `NULL` on error, setting `errno` to `EINVAL`.

The `priv_gettext()` function returns a pointer to a string consisting of one or more newline-separated lines of text describing the privilege. The text is localized using `{LC_MESSAGES}`. The application is responsible for freeing the memory returned.

These functions pick up privileges allocated during the lifetime of the process using `priv_getbyname(9F)` by refreshing the internal data structures when necessary.

RETURN VALUES

Upon successful completion, `priv_str_to_set()` and `priv_set_to_str()` return a non-null pointer to allocated memory that should be freed by the application using the appropriate functions when it is no longer referenced.

The `priv_getbynum()` and `priv_getsetbynum()` functions return non-null pointers to constant memory that should not be modified or freed by the application. Otherwise, `NULL` is returned and `errno` is set to indicate the error.

Upon successful completion, `priv_getbyname()` and `priv_getsetbyname()` return a non-negative integer. Otherwise, -1 is returned and `errno` is set to indicate the error.

Upon successful completion, `priv_gettext()` returns a non-null value. It returns `NULL` if an error occurs or no descriptive text for the specified privilege can be found.

ERRORS

The `priv_str_to_set()` and `priv_set_to_str()` functions will fail if:

<code>ENOMEM</code>	The physical limits of the system are exceeded by the memory allocation needed to hold a privilege set.
<code>EAGAIN</code>	There is not enough memory available to allocate sufficient memory to hold a privilege set, but the application could try again later.

All of these functions will fail if:

<code>EINVAL</code>	One or more of the arguments is invalid.
---------------------	--

EXAMPLES

EXAMPLE 1 List all the sets and privileges defined in the system.

The following example lists all the sets and privileges defined in the system.

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

priv_str_to_set(3C)

EXAMPLE 1 List all the sets and privileges defined in the system. *(Continued)*

```
/* list all the sets and privileges defined in the system */  
  
const char *name;  
int i;  
  
printf("Each process has the following privilege sets:\n");  
for (i = 0; (name = priv_getsetbynum(i++)) != NULL; )  
    printf("\t%s\n", name);  
  
printf("Each set can contain the following privileges:\n");  
for (i = 0; (name = priv_getbynum(i++)) != NULL; )  
    printf("\t%s\n", name);
```

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving
MT-Level	MT-Safe

SEE ALSO [free\(3C\)](#), [priv_set\(3C\)](#), [attributes\(5\)](#), [privileges\(5\)](#), [priv_getbyname\(9F\)](#)

NAME	pset_getloadavg – get system load averages for a processor set						
SYNOPSIS	<pre>#include <sys/pset.h> #include <sys/loadavg.h> int pset_getloadavg(psetid_t pset, double loadavg[], int nelem);</pre>						
DESCRIPTION	<p>The <code>pset_getloadavg()</code> function returns the number of processes assigned to the specified processor set that are in the system run queue, averaged over various periods of time. Up to <i>nelem</i> samples are retrieved and assigned to successive elements of <i>loadavg[]</i>. The system imposes a maximum of 3 samples, representing averages over the last 1, 5, and 15 minutes, respectively.</p> <p>The <code>LOADAVG_1MIN</code>, <code>LOADAVG_5MIN</code>, and <code>LOADAVG_15MIN</code> indices, defined in <code><sys/loadavg.h></code>, can be used to extract the data from the appropriate element of the <i>loadavg[]</i> array.</p> <p>If <i>pset</i> is <code>PS_NONE</code>, the load average for processes not assigned to a processor set is returned.</p> <p>If <i>pset</i> is <code>PS_MYID</code>, the load average for the processor set to which the caller is bound is returned. If the caller is not bound to a processor set, the result is the same as if <code>PS_NONE</code> was specified.</p>						
RETURN VALUES	Upon successful completion, the number of samples actually retrieved is returned. If the load average was unobtainable or the processor set does not exist, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.						
ERRORS	<p>The <code>pset_getloadavg()</code> function will fail if:</p> <p><code>EINVAL</code> The number of elements specified is less than 0, or an invalid processor set ID was specified.</p> <p> The caller is in a non-global zone, the pools facility is active, and the specified processor set is not that of the zone's pool.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Stable</td> </tr> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Stable	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Stable						
MT-Level	Async-Signal-Safe						
SEE ALSO	<code>uptime(1)</code> , <code>w(1)</code> , <code>psrset(1M)</code> , <code>prstat(1M)</code> , <code>pset_bind(2)</code> , <code>pset_create(2)</code> , <code>kstat(3PERL)</code> , <code>attributes(5)</code>						

psignal(3C)

NAME	psignal, psignal – system signal messages				
SYNOPSIS	<pre>#include <siginfo.h> void psignal(int <i>sig</i>, const char *<i>s</i>); void psiginfo(siginfo_t *<i>pinfo</i>, char *<i>s</i>);</pre>				
DESCRIPTION	The <code>psignal()</code> and <code>psiginfo()</code> functions produce messages on the standard error output describing a signal. The <i>sig</i> argument is a signal that may have been passed as the first argument to a signal handler. The <i>pinfo</i> argument is a pointer to a <code>siginfo</code> structure that may have been passed as the second argument to an enhanced signal handler. See <code>sigaction(2)</code> . The argument string <i>s</i> is printed first, followed by a colon and a blank, followed by the message and a NEWLINE character.				
USAGE	If the application is linked with <code>-lintl</code> , then messages printed from these functions are in the native language specified by the <code>LC_MESSAGES</code> locale category. See <code>setlocale(3C)</code> .				
ATTRIBUTES	See <code>attributes(5)</code> 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>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe				
SEE ALSO	<code>sigaction(2)</code> , <code>gettext(3C)</code> , <code>perror(3C)</code> , <code>setlocale(3C)</code> , <code>siginfo.h(3HEAD)</code> , <code>signal.h(3HEAD)</code> , <code>attributes(5)</code>				

NAME	psignal, sys_siglist – system signal messages
SYNOPSIS	<pre> /usr/ucb/cc [flag ...] file ... void psignal(sig, s); unsigned sig; char *s; char *sys_siglist[]; </pre>
DESCRIPTION	<p>psignal() produces a short message on the standard error file describing the indicated signal. First the argument string <i>s</i> is printed, then a colon, then the name of the signal and a NEWLINE. Most usefully, the argument string is the name of the program which incurred the signal. The signal number should be from among those found in <signal.h>.</p> <p>To simplify variant formatting of signal names, the vector of message strings <i>sys_siglist</i> is provided; the signal number can be used as an index in this table to get the signal name without the newline. The define <i>NSIG</i> defined in <signal.h> is the number of messages provided for in the table; it should be checked because new signals may be added to the system before they are added to the table.</p>
SEE ALSO	perror(3C), signal(3C)
NOTES	Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.

pthread_atfork(3C)

NAME	pthread_atfork – register fork handlers		
SYNOPSIS	<pre>#include <sys/types.h> #include <unistd.h> int pthread_atfork(void (*prepare) (void), void (*parent) (void), void (*child) (void));</pre>		
DESCRIPTION	<p>The <code>pthread_atfork()</code> function declares fork handlers to be called prior to and following <code>fork(2)</code>, within the thread that called <code>fork()</code>. The order of calls to <code>pthread_atfork()</code> is significant.</p> <p>Before <code>fork()</code> processing begins, the <i>prepare</i> fork handler is called. The <i>prepare</i> handler is not called if its address is <code>NULL</code>.</p> <p>The <i>parent</i> fork handler is called after <code>fork()</code> processing finishes in the parent process, and the <i>child</i> fork handler is called after <code>fork()</code> processing finishes in the child process. If the address of <i>parent</i> or <i>child</i> is <code>NULL</code>, then its handler is not called.</p> <p>The <i>prepare</i> fork handler is called in LIFO (last-in first-out) order, whereas the <i>parent</i> and <i>child</i> fork handlers are called in FIFO (first-in first-out) order. This calling order allows applications to preserve locking order.</p>		
RETURN VALUES	Upon successful completion, <code>pthread_atfork()</code> returns 0. Otherwise, an error number is returned.		
ERRORS	The <code>pthread_atfork()</code> function will fail if: <table><tr><td>ENOMEM</td><td>Insufficient table space exists to record the fork handler addresses.</td></tr></table>	ENOMEM	Insufficient table space exists to record the fork handler addresses.
ENOMEM	Insufficient table space exists to record the fork handler addresses.		
USAGE	Solaris threads do not offer <code>pthread_atfork()</code> functionality (there is no <code>thr_atfork()</code> interface). However, a Solaris threads application can call <code>pthread_atfork()</code> to ensure <code>fork()</code> -safety, since the two thread APIs are interoperable. See <code>fork(2)</code> for information relating to <code>fork()</code> in a Solaris threads environment in Solaris 10 relative to previous releases.		
EXAMPLES	<p>EXAMPLE 1 mMake a library safe with respect to <code>fork()</code>.</p> <p>All multithreaded applications that call <code>fork()</code> in a POSIX threads program and do more than simply call <code>exec(2)</code> in the child of the fork need to ensure that the child is protected from deadlock.</p> <p>Since the "fork-one" model results in duplicating only the thread that called <code>fork()</code>, it is possible that at the time of the call another thread in the parent owns a lock. This thread is not duplicated in the child, so no thread will unlock this lock in the child. Deadlock occurs if the single thread in the child needs this lock.</p>		

EXAMPLE 1 mMake a library safe with respect to `fork()`. (Continued)

The problem is more serious with locks in libraries. Since a library writer does not know if the application using the library calls `fork()`, the library must protect itself from such a deadlock scenario. If the application that links with this library calls `fork()` and does not call `exec()` in the child, and if it needs a library lock that may be held by some other thread in the parent that is inside the library at the time of the fork, the application deadlocks inside the library.

The following describes how to make a library safe with respect to `fork()` by using `pthread_atfork()`.

1. Identify all locks used by the library (for example $\{L1, \dots, Ln\}$). Identify also the locking order for these locks (for example $\{L1 \dots Ln\}$, as well.)
2. Add a call to `pthread_atfork(f1, f2, f3)` in the library's `.init` section. `f1`, `f2`, `f3` are defined as follows:

```
f1 ( )
{
    /* ordered in lock order */
    pthread_mutex_lock(L1);
    pthread_mutex_lock( . . . );
    pthread_mutex_lock(Ln);
}

f2 ( )
{
    pthread_mutex_unlock(L1);
    pthread_mutex_unlock( . . . );
    pthread_mutex_unlock(Ln);
}

f3 ( )
{
    pthread_mutex_unlock(L1);
    pthread_mutex_unlock( . . . );
    pthread_mutex_unlock(Ln);
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `exec(2)`, `fork(2)`, `atexit(3C)`, `attributes(5)`, `standards(5)`

pthread_attr_getdetachstate(3C)

NAME pthread_attr_getdetachstate, pthread_attr_setdetachstate – get or set detachstate attribute

SYNOPSIS

```
cc -mt [ flag... ] file... -lpthread [ -lrt library... ]
#include <pthread.h>

int pthread_attr_setdetachstate(pthread_attr_t *attr, int detachstate);

int pthread_attr_getdetachstate(const pthread_attr_t *attr, int *detachstate);
```

DESCRIPTION The *detachstate* attribute controls whether the thread is created in a detached state. If the thread is created detached, then use of the ID of the newly created thread by the pthread_detach() or pthread_join() function is an error.

The pthread_attr_setdetachstate() and pthread_attr_getdetachstate(), respectively, set and get the *detachstate* attribute in the *attr* object.

The *detachstate* can be set to either PTHREAD_CREATE_DETACHED or PTHREAD_CREATE_JOINABLE. A value of PTHREAD_CREATE_DETACHED causes all threads created with *attr* to be in the detached state, whereas using a value of PTHREAD_CREATE_JOINABLE causes all threads created with *attr* to be in the joinable state. The default value of the *detachstate* attribute is PTHREAD_CREATE_JOINABLE.

RETURN VALUES Upon successful completion, pthread_attr_setdetachstate() and pthread_attr_getdetachstate() return a value of 0. Otherwise, an error number is returned to indicate the error.

The pthread_attr_getdetachstate() function stores the value of the *detachstate* attribute in *detachstate* if successful.

ERRORS The pthread_attr_setdetachstate() or pthread_attr_getdetachstate() functions may fail if:

EINVAL *attr* or *detachstate* is invalid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO pthread_attr_init(3C), pthread_attr_setstackaddr(3C), pthread_attr_setstacksize(3C), pthread_create(3C), attributes(5), standards(5)

NAME	pthread_attr_getguardsize, pthread_attr_setguardsize – get or set the thread guardsize attribute
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_attr_getguardsize(const pthread_attr_t *restrict attr, size_t *restrict guardsize); int pthread_attr_setguardsize(pthread_attr_t *attr, size_t guardsize);</pre>
DESCRIPTION	<p>The <i>guardsize</i> attribute controls the size of the guard area for the created thread's stack. The <i>guardsize</i> attribute provides protection against overflow of the stack pointer. If a thread's stack is created with guard protection, the implementation allocates extra memory at the overflow end of the stack as a buffer against stack overflow of the stack pointer. If an application overflows into this buffer an error results (possibly in a SIGSEGV signal being delivered to the thread).</p> <p>The <i>guardsize</i> attribute is provided to the application for two reasons:</p> <ol style="list-style-type: none"> 1. Overflow protection can potentially result in wasted system resources. An application that creates a large number of threads, and which knows its threads will never overflow their stack, can save system resources by turning off guard areas. 2. When threads allocate large data structures on the stack, large guard areas may be needed to detect stack overflow. <p>The <code>pthread_attr_getguardsize()</code> function gets the <i>guardsize</i> attribute in the <i>attr</i> object. This attribute is returned in the <i>guardsize</i> parameter.</p> <p>The <code>pthread_attr_setguardsize()</code> function sets the <i>guardsize</i> attribute in the <i>attr</i> object. The new value of this attribute is obtained from the <i>guardsize</i> parameter. If <i>guardsize</i> is 0, a guard area will not be provided for threads created with <i>attr</i>. If <i>guardsize</i> is greater than 0, a guard area of at least size <i>guardsize</i> bytes is provided for each thread created with <i>attr</i>.</p> <p>A conforming implementation is permitted to round up the value contained in <i>guardsize</i> to a multiple of the configurable system variable <code>PAGESIZE</code>. If an implementation rounds up the value of <i>guardsize</i> to a multiple of <code>PAGESIZE</code>, a call to <code>pthread_attr_getguardsize()</code> specifying <i>attr</i> will store in the <i>guardsize</i> parameter the guard size specified by the previous <code>pthread_attr_setguardsize()</code> function call.</p> <p>The default value of the <i>guardsize</i> attribute is <code>PAGESIZE</code> bytes. The actual value of <code>PAGESIZE</code> is implementation-dependent and may not be the same on all implementations.</p> <p>If the <i>stackaddr</i> attribute has been set (that is, the caller is allocating and managing its own thread stacks), the <i>guardsize</i> attribute is ignored and no protection will be provided by the implementation. It is the responsibility of the application to manage stack overflow along with stack allocation and management in this case.</p>

pthread_attr_getguardsize(3C)

RETURN VALUES If successful, the `pthread_attr_getguardsize()` and `pthread_attr_setguardsize()` functions return 0. Otherwise, an error number is returned to indicate the error.

ERRORS The `pthread_attr_getguardsize()` and `pthread_attr_setguardsize()` functions will fail if:

- EINVAL The attribute *attr* is invalid.
- EINVAL The parameter *guardsize* is invalid.
- EINVAL The parameter *guardsize* contains an invalid value.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `sysconf(3C)`, `pthread_attr_init(3C)`, `attributes(5)`, `standards(5)`

NAME	pthread_attr_getinheritsched, pthread_attr_setinheritsched – get or set inheritsched attribute				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_attr_getinheritsched(const pthread_attr_t *restrict <i>attr</i>, int *restrict <i>inheritsched</i>); int pthread_attr_setinheritsched(pthread_attr_t *<i>attr</i>, int <i>inheritsched</i>);</pre>				
DESCRIPTION	<p>The functions pthread_attr_setinheritsched() and pthread_attr_getinheritsched(), respectively, set and get the <i>inheritsched</i> attribute in the <i>attr</i> argument.</p> <p>When the attribute objects are used by pthread_create(), the <i>inheritsched</i> attribute determines how the other scheduling attributes of the created thread are to be set:</p> <table border="0"> <tr> <td style="vertical-align: top;">PTHREAD_INHERIT_SCHED</td> <td>Specifies that the scheduling policy and associated attributes are to be inherited from the creating thread, and the scheduling attributes in this <i>attr</i> argument are to be ignored.</td> </tr> <tr> <td style="vertical-align: top;">PTHREAD_EXPLICIT_SCHED</td> <td>Specifies that the scheduling policy and associated attributes are to be set to the corresponding values from this attribute object.</td> </tr> </table> <p>The symbols PTHREAD_INHERIT_SCHED and PTHREAD_EXPLICIT_SCHED are defined in the header <pthread.h>.</p>	PTHREAD_INHERIT_SCHED	Specifies that the scheduling policy and associated attributes are to be inherited from the creating thread, and the scheduling attributes in this <i>attr</i> argument are to be ignored.	PTHREAD_EXPLICIT_SCHED	Specifies that the scheduling policy and associated attributes are to be set to the corresponding values from this attribute object.
PTHREAD_INHERIT_SCHED	Specifies that the scheduling policy and associated attributes are to be inherited from the creating thread, and the scheduling attributes in this <i>attr</i> argument are to be ignored.				
PTHREAD_EXPLICIT_SCHED	Specifies that the scheduling policy and associated attributes are to be set to the corresponding values from this attribute object.				
RETURN VALUES	If successful, the pthread_attr_setinheritsched() and pthread_attr_getinheritsched() functions return 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>The pthread_attr_setinheritsched() or pthread_attr_getinheritsched() functions may fail if:</p> <p>EINVAL <i>attr</i> or <i>inheritsched</i> is invalid.</p>				
USAGE	After these attributes have been set, a thread can be created with the specified attributes using pthread_create(). Using these routines does not affect the current running thread.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

pthread_attr_getinheritsched(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO pthread_attr_getschedparam(3C), pthread_attr_init(3C), pthread_attr_setscope(3C), pthread_attr_setschedpolicy(3C), pthread_create(3C), pthread_setschedparam(3C), attributes(5), standards(5)

NAME	pthread_attr_getschedparam, pthread_attr_setschedparam – get or set schedparam attribute						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_attr_getschedparam(const pthread_attr_t *restrict attr, struct sched_param *restrict param); int pthread_attr_setschedparam(pthread_attr_t *restrict attr, const struct sched_param *restrict param);</pre>						
DESCRIPTION	The functions pthread_attr_setschedparam() and pthread_attr_getschedparam(), respectively, set and get the scheduling parameter attributes in the <i>attr</i> argument. The contents of the <i>param</i> structure are defined in <sched.h>. For the SCHED_FIFO and SCHED_RR policies, the only required member of <i>param</i> is <i>sched_priority</i> .						
RETURN VALUES	If successful, the pthread_attr_setschedparam() and pthread_attr_getschedparam() functions return 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	<p>The pthread_attr_setschedparam() function may fail if:</p> <p>EINVAL <i>attr</i> is invalid.</p> <p>The pthread_attr_getschedparam() function may fail if:</p> <p>EINVAL <i>attr</i> or <i>param</i> is invalid.</p>						
USAGE	After these attributes have been set, a thread can be created with the specified attributes using pthread_create(). Using these routines does not affect the current running thread.						
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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	pthread_attr_init(3C), pthread_attr_setscope(3C), pthread_attr_setinheritsched(3C), pthread_attr_setschedpolicy(3C), pthread_create(3C), pthread_setschedparam(3C), attributes(5), standards(5)						

pthread_attr_getschedpolicy(3C)

NAME	pthread_attr_getschedpolicy, pthread_attr_setschedpolicy – get or set schedpolicy attribute						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_attr_getschedpolicy(const pthread_attr_t *restrict <i>attr</i>, int *restrict <i>policy</i>); int pthread_attr_setschedpolicy(pthread_attr_t *<i>attr</i>, int <i>policy</i>);</pre>						
DESCRIPTION	<p>The functions pthread_attr_setschedpolicy() and pthread_attr_getschedpolicy(), respectively, set and get the <i>schedpolicy</i> attribute in the <i>attr</i> argument.</p> <p>The supported values of <i>policy</i> include SCHED_FIFO, SCHED_RR and SCHED_OTHER, which are defined by the header <sched.h>. When threads executing with the scheduling policy SCHED_FIFO or SCHED_RR are waiting on a mutex, they acquire the mutex in priority order when the mutex is unlocked.</p>						
RETURN VALUES	If successful, the pthread_attr_setschedpolicy() and pthread_attr_getschedpolicy() functions return 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	The pthread_attr_setschedpolicy() or pthread_attr_getschedpolicy() function may fail if:						
	EINVAL <i>attr</i> or <i>policy</i> is invalid.						
USAGE	After these attributes have been set, a thread can be created with the specified attributes using pthread_create(). Using these routines does not affect the current running thread.						
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>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	pthread_attr_init(3C), pthread_attr_setscope(3C), pthread_attr_setinheritsched(3C), pthread_attr_setschedparam(3C), pthread_create(3C), pthread_setschedparam(3C), attributes(5), standards(5)						

NAME	pthread_attr_getscope, pthread_attr_setscope – get or set contentionscope attribute						
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_attr_getscope(const pthread_attr_t *restrict attr, int *restrict contentionscope); int pthread_attr_setscope(pthread_attr_t *attr, int contentionscope);</pre>						
DESCRIPTION	<p>The pthread_attr_setscope() and pthread_attr_getscope() functions are used to set and get the <i>contentionscope</i> attribute in the <i>attr</i> object.</p> <p>The <i>contentionscope</i> attribute can have the value PTHREAD_SCOPE_SYSTEM, signifying system scheduling contention scope, or PTHREAD_SCOPE_PROCESS, signifying process scheduling contention scope.</p> <p>The symbols PTHREAD_SCOPE_SYSTEM and PTHREAD_SCOPE_PROCESS are defined by the header <pthread.h>.</p>						
RETURN VALUES	If successful, the pthread_attr_setscope() and pthread_attr_getscope() functions return 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	<p>The pthread_attr_setscope(), or pthread_attr_getscope(), function may fail if:</p> <p>EINVAL <i>attr</i> or <i>contentionscope</i> is invalid.</p>						
USAGE	After these attributes have been set, a thread can be created with the specified attributes using pthread_create(). Using these routines does not affect the current running thread.						
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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	pthread_attr_init(3C), pthread_attr_setinheritsched(3C), pthread_attr_setschedpolicy(3C), pthread_attr_setschedparam(3C), pthread_create(3C), pthread_setschedparam(3C), attributes(5), standards(5)						

pthread_attr_getstack(3C)

NAME	pthread_attr_getstack, pthread_attr_setstack – get and set stack attributes						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_attr_getstack(const pthread_attr_t *restrict attr, void **restrict stackaddr, size_t *restrict stacksize); int pthread_attr_setstack(pthread_attr_t * attr, void *stackaddr, size_t stacksize);</pre>						
DESCRIPTION	<p>The pthread_attr_getstack() and pthread_attr_setstack() functions, respectively, get and set the thread creation stack attributes <i>stackaddr</i> and <i>stacksize</i> in the <i>attr</i> object.</p> <p>The stack attributes specify the area of storage to be used for the created thread's stack. The base (lowest addressable byte) of the storage is <i>stackaddr</i>, and the size of the storage is <i>stacksize</i> bytes. The <i>stacksize</i> argument must be at least {PTHREAD_STACK_MIN}. The <i>stackaddr</i> argument must be aligned appropriately to be used as a stack; for example, pthread_attr_setstack() might fail with EINVAL if (<i>stackaddr</i> & 0x7) is not 0. All pages within the stack described by <i>stackaddr</i> and <i>stacksize</i> are both readable and writable by the thread.</p>						
RETURN VALUES	<p>Upon successful completion, these functions return a 0; otherwise, an error number is returned to indicate the error.</p> <p>The pthread_attr_getstack() function stores the stack attribute values in <i>stackaddr</i> and <i>stacksize</i> if successful.</p>						
ERRORS	<p>The pthread_attr_setstack() function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EINVAL</td> <td>The value of <i>stacksize</i> is less than {PTHREAD_STACK_MIN}.</td> </tr> </table> <p>The pthread_attr_setstack() function may fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EACCES</td> <td>The stack page(s) described by <i>stackaddr</i> and <i>stacksize</i> are not both readable and writable by the thread.</td> </tr> <tr> <td style="padding-right: 20px;">EINVAL</td> <td>The value of <i>stackaddr</i> does not have proper alignment to be used as a stack, or (<i>stackaddr</i> + <i>stacksize</i>) lacks proper alignment.</td> </tr> </table>	EINVAL	The value of <i>stacksize</i> is less than {PTHREAD_STACK_MIN}.	EACCES	The stack page(s) described by <i>stackaddr</i> and <i>stacksize</i> are not both readable and writable by the thread.	EINVAL	The value of <i>stackaddr</i> does not have proper alignment to be used as a stack, or (<i>stackaddr</i> + <i>stacksize</i>) lacks proper alignment.
EINVAL	The value of <i>stacksize</i> is less than {PTHREAD_STACK_MIN}.						
EACCES	The stack page(s) described by <i>stackaddr</i> and <i>stacksize</i> are not both readable and writable by the thread.						
EINVAL	The value of <i>stackaddr</i> does not have proper alignment to be used as a stack, or (<i>stackaddr</i> + <i>stacksize</i>) lacks proper alignment.						
USAGE	These functions are appropriate for use by applications in an environment where the stack for a thread must be placed in some particular region of memory.						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:						

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

`pthread_attr_getstack(3C)`

SEE ALSO `pthread_attr_init(3C)`, `pthread_attr_setdetachstate(3C)`,
`pthread_attr_setstacksize(3C)`, `pthread_create(3C)`, `attributes(5)`

pthread_attr_getstackaddr(3C)

NAME pthread_attr_getstackaddr, pthread_attr_setstackaddr – get or set stackaddr attribute

SYNOPSIS

```
cc -mt [ flag... ] file... -lpthread [ -lrt library... ]
#include <pthread.h>

int pthread_attr_getstackaddr(const pthread_attr_t *restrict attr,
    void **restrict stackaddr);

int pthread_attr_setstackaddr(pthread_attr_t *attr, void *stackaddr);
```

DESCRIPTION

The functions pthread_attr_setstackaddr() and pthread_attr_getstackaddr(), respectively, set and get the thread creation stackaddr attribute in the attr object. The stackaddr default is NULL. See pthread_create(3C).

The stackaddr attribute specifies the location of storage to be used for the created thread's stack. The size of the storage is at least PTHREAD_STACK_MIN.

RETURN VALUES

Upon successful completion, pthread_attr_setstackaddr() and pthread_attr_getstackaddr() return a value of 0. Otherwise, an error number is returned to indicate the error.

If successful, the pthread_attr_getstackaddr() function stores the stackaddr attribute value in stackaddr.

ERRORS

The pthread_attr_setstackaddr() function may fail if:

EINVAL attr is invalid.

The pthread_attr_getstackaddr() function may fail if:

EINVAL attr or stackaddr is invalid.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO pthread_attr_init(3C), pthread_attr_setdetachstate(3C), pthread_attr_setstacksize(3C), pthread_create(3C), attributes(5), standards(5)

NAME	pthread_attr_getstacksize, pthread_attr_setstacksize – get or set stacksize attribute						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_attr_getstacksize(const pthread_attr_t *restrict attr, size_t *restrict stacksize); int pthread_attr_setstacksize(pthread_attr_t *attr, size_t stacksize);</pre>						
DESCRIPTION	<p>The functions pthread_attr_setstacksize() and pthread_attr_getstacksize(), respectively, set and get the thread creation stacksize attribute in the attr object.</p> <p>The stacksize attribute defines the minimum stack size (in bytes) allocated for the created threads stack. When the stacksize argument is NULL, the default stack size becomes 1 megabyte for 32-bit processes and 2 megabytes for 64-bit processes.</p>						
RETURN VALUES	Upon successful completion, pthread_attr_setstacksize() and pthread_attr_getstacksize() return a value of 0. Otherwise, an error number is returned to indicate the error. The pthread_attr_getstacksize() function stores the stacksize attribute value in stacksize if successful.						
ERRORS	<p>The pthread_attr_setstacksize() or pthread_attr_getstacksize() function may fail if:</p> <p>EINVAL <i>attr</i> or <i>stacksize</i> 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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	pthread_attr_init(3C), pthread_attr_setstackaddr(3C), pthread_attr_setdetachstate(3C), pthread_create(3C), attributes(5), standards(5)						

pthread_attr_init(3C)

NAME	pthread_attr_init, pthread_attr_destroy – initialize or destroy threads attribute object																											
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_attr_init(pthread_attr_t *attr); int pthread_attr_destroy(pthread_attr_t *attr);</pre>																											
DESCRIPTION	<p>The function <code>pthread_attr_init()</code> initializes a thread attributes object <i>attr</i> with the default value for all of the individual attributes used by a given implementation.</p> <p>The resulting attribute object (possibly modified by setting individual attribute values), when used by <code>pthread_create()</code>, defines the attributes of the thread created. A single attributes object can be used in multiple simultaneous calls to <code>pthread_create()</code>.</p> <p>The <code>pthread_attr_init()</code> function initializes a thread attributes object (<i>attr</i>) with the default value for each attribute as follows:</p> <table><thead><tr><th>Attribute</th><th>Default Value</th><th>Meaning of Default</th></tr></thead><tbody><tr><td><i>contentionscope</i></td><td>PTHREAD_SCOPE_PROCESS</td><td>resource competition within process</td></tr><tr><td><i>detachstate</i></td><td>PTHREAD_CREATE_JOINABLE</td><td>joinable by other threads</td></tr><tr><td><i>stackaddr</i></td><td>NULL</td><td>stack allocated by system</td></tr><tr><td><i>stacksize</i></td><td>NULL</td><td>1 or 2 megabyte</td></tr><tr><td><i>priority</i></td><td>0</td><td>priority of the thread</td></tr><tr><td><i>policy</i></td><td>SCHED_OTHER</td><td>determined by system</td></tr><tr><td><i>inheritsched</i></td><td>PTHREAD_EXPLICIT_SCHED</td><td>scheduling policy and parameters not inherited but explicitly defined by the attribute object</td></tr><tr><td><i>guardsize</i></td><td>PAGESIZE</td><td>size of guard area for a thread's created stack</td></tr></tbody></table> <p>The <code>pthread_attr_destroy()</code> function destroys a thread attributes object (<i>attr</i>), which cannot be reused until it is reinitialized. An implementation may cause <code>pthread_attr_destroy()</code> to set <i>attr</i> to an implementation-dependent invalid value. The behavior of using the attribute after it has been destroyed is undefined.</p>	Attribute	Default Value	Meaning of Default	<i>contentionscope</i>	PTHREAD_SCOPE_PROCESS	resource competition within process	<i>detachstate</i>	PTHREAD_CREATE_JOINABLE	joinable by other threads	<i>stackaddr</i>	NULL	stack allocated by system	<i>stacksize</i>	NULL	1 or 2 megabyte	<i>priority</i>	0	priority of the thread	<i>policy</i>	SCHED_OTHER	determined by system	<i>inheritsched</i>	PTHREAD_EXPLICIT_SCHED	scheduling policy and parameters not inherited but explicitly defined by the attribute object	<i>guardsize</i>	PAGESIZE	size of guard area for a thread's created stack
Attribute	Default Value	Meaning of Default																										
<i>contentionscope</i>	PTHREAD_SCOPE_PROCESS	resource competition within process																										
<i>detachstate</i>	PTHREAD_CREATE_JOINABLE	joinable by other threads																										
<i>stackaddr</i>	NULL	stack allocated by system																										
<i>stacksize</i>	NULL	1 or 2 megabyte																										
<i>priority</i>	0	priority of the thread																										
<i>policy</i>	SCHED_OTHER	determined by system																										
<i>inheritsched</i>	PTHREAD_EXPLICIT_SCHED	scheduling policy and parameters not inherited but explicitly defined by the attribute object																										
<i>guardsize</i>	PAGESIZE	size of guard area for a thread's created stack																										
RETURN VALUES	Upon successful completion, <code>pthread_attr_init()</code> and <code>pthread_attr_destroy()</code> return a value of 0. Otherwise, an error number is returned to indicate the error.																											
ERRORS	The <code>pthread_attr_init()</code> function will fail if:																											

pthread_attr_init(3C)

ENOMEM Insufficient memory exists to initialize the thread attributes object.

The pthread_attr_destroy() function may fail if:

EINVAL *attr* is invalid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO sysconf(3C), pthread_attr_getdetachstate(3C),
pthread_attr_getguardsize(3C), pthread_attr_getinheritsched(3C),
pthread_attr_getschedparam(3C), pthread_attr_getschedpolicy(3C),
pthread_attr_getscope(3C), pthread_attr_getstackaddr(3C),
pthread_attr_getstacksize(3C), pthread_attr_setdetachstate(3C),
pthread_attr_setguardsize(3C), pthread_attr_setinheritsched(3C),
pthread_attr_setschedparam(3C), pthread_attr_setschedpolicy(3C),
pthread_attr_setscope(3C), pthread_attr_setstackaddr(3C),
pthread_attr_setstacksize(3C), pthread_create(3C), attributes(5),
standards(5)

pthread_barrierattr_destroy(3C)

NAME	pthread_barrierattr_destroy, pthread_barrierattr_init – destroy and initialize barrier attributes object						
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <pthread.h> int pthread_barrierattr_destroy(pthread_barrierattr_t *attr); int pthread_barrierattr_init(pthread_barrierattr_t *attr);</pre>						
DESCRIPTION	<p>The pthread_barrierattr_destroy() function destroys a barrier attributes object. A destroyed <i>attr</i> attributes object can be reinitialized using pthread_barrierattr_init(). The results of otherwise referencing the object after it has been destroyed are undefined. An implementation can cause pthread_barrierattr_destroy() to set the object referenced by <i>attr</i> to an invalid value.</p> <p>The pthread_barrierattr_init() function initializes a barrier attributes object <i>attr</i> with the default value for all of the attributes defined by the implementation.</p> <p>Results are undefined if pthread_barrierattr_init() is called specifying an already initialized <i>attr</i> attributes object.</p> <p>After a barrier attributes object has been used to initialize one or more barriers, any function affecting the attributes object (including destruction) does not affect any previously initialized barrier.</p>						
RETURN VALUES	Upon successful completion, the pthread_barrierattr_destroy() and pthread_barrierattr_init() functions returns 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	The pthread_barrierattr_init() function will fail if: ENOMEM Insufficient memory exists to initialize the barrier attributes object. The pthread_barrierattr_destroy() function may fail if: EINVAL The value specified by <i>attr</i> is invalid.						
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>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	pthread_barrierattr_getpshared(3C), attributes(5), standards(5)						

NAME	pthread_barrierattr_getpshared, pthread_barrierattr_setpshared – get and set process-shared attribute of barrier attributes object
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <pthread.h> int pthread_barrierattr_getpshared(const pthread_barrierattr_t *restrict attr, int *restrict pshared); int pthread_barrierattr_setpshared(pthread_barrierattr_t *attr, int pshared);</pre>
DESCRIPTION	<p>The pthread_barrierattr_getpshared() function obtains the value of the <i>process-shared</i> attribute from the attributes object referenced by <i>attr</i>. The pthread_barrierattr_setpshared() function sets the <i>process-shared</i> attribute in an initialized attributes object referenced by <i>attr</i>.</p> <p>The <i>process-shared</i> attribute is set to PTHREAD_PROCESS_SHARED to permit a barrier to be operated upon by any thread that has access to the memory where the barrier is allocated. If the <i>process-shared</i> attribute is PTHREAD_PROCESS_PRIVATE, the barrier will only be operated upon by threads created within the same process as the thread that initialized the barrier. If threads of different processes attempt to operate on such a barrier, the behavior is undefined.</p> <p>The default value of the attribute is PTHREAD_PROCESS_PRIVATE. Both constants PTHREAD_PROCESS_SHARED and PTHREAD_PROCESS_PRIVATE are defined in <pthread.h>.</p> <p>No barrier attributes other than the <i>process-shared</i> attribute are provided.</p>
RETURN VALUES	<p>Upon successful completion, the pthread_barrierattr_getpshared() function returns 0 and stores the value of the <i>process-shared</i> attribute of <i>attr</i> into the object referenced by the <i>pshared</i> parameter. Otherwise, an error number is returned to indicate the error.</p> <p>Upon successful completion, the pthread_barrierattr_setpshared() function returns 0. Otherwise, an error number is returned to indicate the error.</p>
ERRORS	<p>These functions may fail if:</p> <p>EINVAL The value specified by <i>attr</i> is invalid.</p> <p>The pthread_barrierattr_setpshared() function may fail if:</p> <p>EINVAL The new value specified for the <i>process-shared</i> attribute is not one of the legal values PTHREAD_PROCESS_SHARED or PTHREAD_PROCESS_PRIVATE.</p>

pthread_barrierattr_getpshared(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `pthread_barrier_init(3C)`, `pthread_barrierattr_destroy(3C)`, `attributes(5)`, `standards(5)`

NAME	pthread_barrier_destroy, pthread_barrier_init – destroy and initialize a barrier object						
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <pthread.h> int pthread_barrier_destroy(pthread_barrier_t *barrier); int pthread_barrier_init(pthread_barrier_t *barrier, const pthread_barrierattr_t *restrict attr, unsigned count);</pre>						
DESCRIPTION	<p>The <code>pthread_barrier_destroy()</code> function destroys the barrier referenced by <i>barrier</i> and releases any resources used by the barrier. The effect of subsequent use of the barrier is undefined until the barrier is reinitialized by another call to <code>pthread_barrier_init()</code>. An implementation can use this function to set barrier to an invalid value. The results are undefined if <code>pthread_barrier_destroy()</code> is called when any thread is blocked on the barrier, or if this function is called with an uninitialized barrier.</p> <p>The <code>pthread_barrier_init()</code> function allocates any resources required to use the barrier referenced by <i>barrier</i> and initializes the barrier with attributes referenced by <i>attr</i>. If <i>attr</i> is <code>NULL</code>, the default barrier attributes are used; the effect is the same as passing the address of a default barrier attributes object. The results are undefined if <code>pthread_barrier_init()</code> is called when any thread is blocked on the barrier (that is, has not returned from the <code>pthread_barrier_wait(3C)</code> call). The results are undefined if a barrier is used without first being initialized. The results are undefined if <code>pthread_barrier_init()</code> is called specifying an already initialized barrier.</p> <p>The <i>count</i> argument specifies the number of threads that must call <code>pthread_barrier_wait()</code> before any of them successfully return from the call. The value specified by <i>count</i> must be greater than 0.</p> <p>If the <code>pthread_barrier_init()</code> function fails, the barrier is not initialized and the contents of <i>barrier</i> are undefined.</p> <p>Only the object referenced by <i>barrier</i> can be used for performing synchronization. The result of referring to copies of that object in calls to <code>pthread_barrier_destroy()</code> or <code>pthread_barrier_wait()</code> is undefined.</p>						
RETURN VALUES	Upon successful completion, these functions returns 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	<p>The <code>pthread_barrier_init()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EAGAIN</td> <td>The system lacks the necessary resources to initialize another barrier.</td> </tr> <tr> <td>EINVAL</td> <td>The value specified by <i>count</i> is equal to 0.</td> </tr> <tr> <td>ENOMEM</td> <td>Insufficient memory exists to initialize the barrier.</td> </tr> </table> <p>The <code>pthread_barrier_init()</code> function may fail if:</p>	EAGAIN	The system lacks the necessary resources to initialize another barrier.	EINVAL	The value specified by <i>count</i> is equal to 0.	ENOMEM	Insufficient memory exists to initialize the barrier.
EAGAIN	The system lacks the necessary resources to initialize another barrier.						
EINVAL	The value specified by <i>count</i> is equal to 0.						
ENOMEM	Insufficient memory exists to initialize the barrier.						

pthread_barrier_destroy(3C)

EBUSY The implementation has detected an attempt to destroy a barrier while it is in use (for example, while being used in a `pthread_barrier_wait()` call) by another thread.

EINVAL The value specified by *attr* is invalid.

The `pthread_barrier_destroy()` function may fail if:

EBUSY The implementation has detected an attempt to destroy a barrier while it is in use (for example, while being used in a `pthread_barrier_wait()` call) by another thread.

EINVAL The value specified by *barrier* is invalid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `pthread_barrier_wait(3C)`, `attributes(5)`, `standards(5)`

NAME	pthread_barrier_wait – synchronize at a barrier
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> [<i>library...</i>] #include <pthread.h> int pthread_barrier_wait(pthread_barrier_t *<i>barrier</i>);</pre>
DESCRIPTION	<p>The pthread_barrier_wait() function synchronizes participating threads at the barrier referenced by <i>barrier</i>. The calling thread blocks until the required number of threads have called pthread_barrier_wait() specifying the barrier.</p> <p>When the required number of threads have called pthread_barrier_wait() specifying the barrier, the constant PTHREAD_BARRIER_SERIAL_THREAD is returned to one unspecified thread and 0 is returned to each of the remaining threads. At this point, the barrier is reset to the state it had as a result of the most recent pthread_barrier_init(3C) function that referenced it.</p> <p>The constant PTHREAD_BARRIER_SERIAL_THREAD is defined in <pthread.h> and its value is distinct from any other value returned by pthread_barrier_wait().</p> <p>The results are undefined if this function is called with an uninitialized barrier.</p> <p>If a signal is delivered to a thread blocked on a barrier, upon return from the signal handler the thread resumes waiting at the barrier if the barrier wait has not completed (that is, if the required number of threads have not arrived at the barrier during the execution of the signal handler); otherwise, the thread continues as normal from the completed barrier wait. Until the thread in the signal handler returns from it, it is unspecified whether other threads may proceed past the barrier once they have all reached it.</p> <p>A thread that has blocked on a barrier does not prevent any unblocked thread that is eligible to use the same processing resources from eventually making forward progress in its execution.</p> <p>Eligibility for processing resources is determined by the scheduling policy.</p>
RETURN VALUES	Upon successful completion, the pthread_barrier_wait() function returns PTHREAD_BARRIER_SERIAL_THREAD for a single (arbitrary) thread synchronized at the barrier and 0 for each of the other threads. Otherwise, an error number is returned to indicate the error.
ERRORS	<p>The pthread_barrier_wait() function will fail if:</p> <p>EINVAL The value specified by <i>barrier</i> does not refer to an initialized barrier object.</p>
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

pthread_barrier_wait(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [pthread_barrier_destroy\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	pthread_cancel – cancel execution of a thread						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_cancel(pthread_t <i>target_thread</i>);</pre>						
DESCRIPTION	<p>The pthread_cancel() function requests that <i>target_thread</i> be canceled.</p> <p>By default, cancellation is deferred until <i>target_thread</i> reaches a cancellation point. See cancellation(5).</p> <p>Cancellation cleanup handlers for <i>target_thread</i> are called when the cancellation is acted on. Upon return of the last cancellation cleanup handler, the thread-specific data destructor functions are called for <i>target_thread</i>. <i>target_thread</i> is terminated when the last destructor function returns.</p> <p>The cancellation processing in <i>target_thread</i> runs asynchronously with respect to the calling thread returning from pthread_cancel().</p>						
RETURN VALUES	If successful, the pthread_cancel() function returns 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	<p>The pthread_cancel() function may fail if:</p> <p>ESRCH No thread was found with an ID corresponding to that specified by the given thread ID, <i>target_thread</i>.</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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	pthread_cleanup_pop(3C), pthread_cleanup_push(3C), pthread_cond_wait(3C), pthread_cond_timedwait(3C), pthread_exit(3C), pthread_join(3C), pthread_setcancelstate(3C), pthread_setcanceltype(3C), pthread_testcancel(3C), setjmp(3C), attributes(5), cancellation(5), condition(5), standards(5)						
NOTES	See cancellation(5) for a discussion of cancellation concepts.						

pthread_cleanup_pop(3C)

NAME	pthread_cleanup_pop – pop a thread cancellation cleanup handler						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> void pthread_cleanup_pop(int <i>execute</i>);</pre>						
DESCRIPTION	<p>The <code>pthread_cleanup_pop()</code> function removes the cleanup handler routine at the top of the cancellation cleanup stack of the calling thread and executes it if <i>execute</i> is non-zero.</p> <p>When the thread calls <code>pthread_cleanup_pop()</code> with a non-zero <i>execute</i> argument, the argument at the top of the stack is popped and executed. An argument of 0 pops the handler without executing it.</p> <p>The Solaris system generates a compile time error if <code>pthread_cleanup_push()</code> does not have a matching <code>pthread_cleanup_pop()</code>.</p> <p>Using <code>longjmp()</code> or <code>siglongjmp()</code> to jump into or out of a push/pop pair can result in either the matching push or the matching pop statement not getting executed.</p>						
RETURN VALUES	The <code>pthread_cleanup_pop()</code> function returns no value.						
ERRORS	No errors are defined.						
	The <code>pthread_cleanup_pop()</code> function will not return an error code of <code>EINTR</code> .						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>pthread_cancel(3C)</code> , <code>pthread_cleanup_push(3C)</code> , <code>pthread_exit(3C)</code> , <code>pthread_join(3C)</code> , <code>pthread_setcancelstate(3C)</code> , <code>pthread_setcanceltype(3C)</code> , <code>pthread_testcancel(3C)</code> , <code>setjmp(3C)</code> , <code>attributes(5)</code> , <code>cancellation(5)</code> , <code>condition(5)</code> , <code>standards(5)</code>						
NOTES	See <code>cancellation(5)</code> for a discussion of cancellation concepts.						

NAME	pthread_cleanup_push – push a thread cancellation cleanup handler						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> void pthread_cleanup_push(void (*<i>handler</i>, void *), void *<i>arg</i>);</pre>						
DESCRIPTION	<p>The pthread_cleanup_push() function pushes the specified cancellation cleanup handler routine, <i>handler</i>, onto the cancellation cleanup stack of the calling thread.</p> <p>When a thread exits or is canceled and its cancellation cleanup stack is not empty, the cleanup handlers are invoked with the argument <i>arg</i> in last in, first out (LIFO) order from the cancellation cleanup stack.</p> <p>The Solaris system generates a compile time error if pthread_cleanup_push() does not have a matching pthread_cleanup_pop().</p> <p>Using longjmp() or siglongjmp() to jump into or out of a push/pop pair can cause either the matching push or the matching pop statement not getting executed.</p>						
RETURN VALUES	The pthread_cleanup_push() function returns no value.						
ERRORS	<p>No errors are defined.</p> <p>The pthread_cleanup_push() function will not return an error code of EINTR.</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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	longjmp(3C), pthread_cancel(3C), pthread_cleanup_pop(3C), pthread_exit(3C), pthread_join(3C), pthread_setcancelstate(3C), pthread_setcanceltype(3C), pthread_testcancel(3C), attributes(5), cancellation(5), condition(5), standards(5)						
NOTES	See cancellation(5) for a discussion of cancellation concepts.						

pthread_condattr_getclock(3C)

NAME	pthread_condattr_getclock, pthread_condattr_setclock – get and set the clock selection condition variable attribute						
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <pthread.h> int pthread_condattr_getclock(const pthread_condattr_t *restrict attr, clockid_t *restrict clock_id); int pthread_condattr_setclock(pthread_condattr_t *attr, clockid_t clock_id);</pre>						
DESCRIPTION	<p>The pthread_condattr_getclock() function obtains the value of the clock attribute from the attributes object referenced by <i>attr</i>. The pthread_condattr_setclock() function sets the clock attribute in an initialized attributes object referenced by <i>attr</i>. If pthread_condattr_setclock() is called with a <i>clock_id</i> argument that refers to a CPU-time clock, the call fails.</p> <p>The clock attribute is the clock ID of the clock that is used to measure the timeout service of pthread_cond_timedwait(3C). The default value of the clock attribute refers to the system clock.</p>						
RETURN VALUES	<p>Upon successful completion, the pthread_condattr_getclock() function returns 0 and stores the value of the clock attribute of <i>attr</i> into the object referenced by the <i>clock_id</i> argument. Otherwise, an error number is returned to indicate the error.</p> <p>Upon successful completion, the pthread_condattr_setclock() function returns 0. Otherwise, an error number is returned to indicate the error.</p>						
ERRORS	<p>These functions may fail if:</p> <p>EINVAL The value specified by <i>attr</i> is invalid.</p> <p>The pthread_condattr_setclock() function may fail if:</p> <p>EINVAL The value specified by <i>clock_id</i> does not refer to a known clock, or is a CPU-time clock.</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>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	pthread_cond_init(3C), pthread_cond_timedwait(3C), pthread_condattr_destroy(3C), pthread_condattr_getpshared(3C), pthread_create(3C), pthread_mutex_init(3C), attributes(5), standards(5)						

NAME	pthread_condattr_getpshared, pthread_condattr_setpshared – get or set the process-shared condition variable attributes
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_condattr_getpshared(const pthread_condattr_t *restrict <i>attr</i>, int *restrict <i>pshared</i>); int pthread_condattr_setpshared(pthread_condattr_t *<i>attr</i>, int <i>pshared</i>);</pre>
DESCRIPTION	<p>The pthread_condattr_getpshared() function obtains the value of the <i>process-shared</i> attribute from the attributes object referenced by <i>attr</i>. The pthread_condattr_setpshared() function is used to set the <i>process-shared</i> attribute in an initialized attributes object referenced by <i>attr</i>.</p> <p>The <i>process-shared</i> attribute is set to PTHREAD_PROCESS_SHARED to permit a condition variable to be operated upon by any thread that has access to the memory where the condition variable is allocated, even if the condition variable is allocated in memory that is shared by multiple processes. If the <i>process-shared</i> attribute is PTHREAD_PROCESS_PRIVATE, the condition variable will only be operated upon by threads created within the same process as the thread that initialized the condition variable; if threads of differing processes attempt to operate on such a condition variable, the behavior is undefined. The default value of the attribute is PTHREAD_PROCESS_PRIVATE.</p> <p>Additional attributes, their default values, and the names of the associated functions to get and set those attribute values are implementation-dependent.</p>
RETURN VALUES	<p>If successful, the pthread_condattr_setpshared() function returns 0. Otherwise, an error number is returned to indicate the error.</p> <p>If successful, the pthread_condattr_getpshared() function returns 0 and stores the value of the <i>process-shared</i> attribute of <i>attr</i> into the object referenced by the <i>pshared</i> parameter. Otherwise, an error number is returned to indicate the error.</p>
ERRORS	<p>The pthread_condattr_getpshared() and pthread_condattr_setpshared() functions may fail if:</p> <p>EINVAL The value specified by <i>attr</i> is invalid.</p> <p>The pthread_condattr_setpshared() function will fail if:</p> <p>EINVAL The new value specified for the attribute is outside the range of legal values for that attribute.</p>
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

pthread_condattr_getpshared(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [pthread_condattr_init\(3C\)](#), [pthread_create\(3C\)](#),
[pthread_mutex_init\(3C\)](#), [pthread_cond_init\(3C\)](#), [attributes\(5\)](#),
[standards\(5\)](#)

NAME	pthread_condattr_init, pthread_condattr_destroy – initialize or destroy condition variable attributes object				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_condattr_init(pthread_condattr_t *attr); int pthread_condattr_destroy(pthread_condattr_t *attr);</pre>				
DESCRIPTION	<p>The pthread_condattr_init() function initializes a condition variable attributes object <i>attr</i> with the default value for all of the attributes defined by the implementation.</p> <p>At present, the only attribute available is the scope of condition variables. The default scope of the attribute is PTHREAD_PROCESS_PRIVATE.</p> <p>Attempts to initialize previously initialized condition variable attributes object will leave the storage allocated by the previous initialization unallocated.</p> <p>After a condition variable attributes object has been used to initialize one or more condition variables, any function affecting the attributes object (including destruction) does not affect any previously initialized condition variables.</p> <p>The pthread_condattr_destroy() function destroys a condition variable attributes object; the object becomes, in effect, uninitialized. An implementation may cause pthread_condattr_destroy() to set the object referenced by <i>attr</i> to an invalid value. A destroyed condition variable attributes object can be re-initialized using pthread_condattr_init(); the results of otherwise referencing the object after it has been destroyed are undefined.</p> <p>Additional attributes, their default values, and the names of the associated functions to get and set those attribute values are implementation-dependent.</p>				
RETURN VALUES	If successful, the pthread_condattr_init() and pthread_condattr_destroy() functions return 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>The pthread_condattr_init() function will fail if:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">ENOMEM</td> <td>Insufficient memory exists to initialize the condition variable attributes object.</td> </tr> </table> <p>The pthread_condattr_destroy() function may fail if:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">EINVAL</td> <td>The value specified by <i>attr</i> is invalid.</td> </tr> </table>	ENOMEM	Insufficient memory exists to initialize the condition variable attributes object.	EINVAL	The value specified by <i>attr</i> is invalid.
ENOMEM	Insufficient memory exists to initialize the condition variable attributes object.				
EINVAL	The value specified by <i>attr</i> is invalid.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

pthread_condattr_init(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [pthread_condattr_getpshared\(3C\)](#), [pthread_condattr_setpshared\(3C\)](#), [pthread_cond_init\(3C\)](#), [pthread_create\(3C\)](#), [pthread_mutex_init\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	pthread_cond_init, pthread_cond_destroy – initialize or destroy condition variables				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_cond_init(pthread_cond_t *restrict <i>cond</i>, const pthread_condattr_t *restrict <i>attr</i>); int pthread_cond_destroy(pthread_cond_t *<i>cond</i>); pthread_cond_t <i>cond</i>= PTHREAD_COND_INITIALIZER;</pre>				
DESCRIPTION	<p>The function <code>pthread_cond_init()</code> initializes the condition variable referenced by <i>cond</i> with attributes referenced by <i>attr</i>. If <i>attr</i> is NULL, the default condition variable attributes are used; the effect is the same as passing the address of a default condition variable attributes object. See <code>pthread_condattr_init(3C)</code>. Upon successful initialization, the state of the condition variable becomes initialized.</p> <p>Attempting to initialize an already initialized condition variable results in undefined behavior.</p> <p>The function <code>pthread_cond_destroy()</code> destroys the given condition variable specified by <i>cond</i>; the object becomes, in effect, uninitialized. An implementation may cause <code>pthread_cond_destroy()</code> to set the object referenced by <i>cond</i> to an invalid value. A destroyed condition variable object can be re-initialized using <code>pthread_cond_init()</code>; the results of otherwise referencing the object after it has been destroyed are undefined.</p> <p>It is safe to destroy an initialized condition variable upon which no threads are currently blocked. Attempting to destroy a condition variable upon which other threads are currently blocked results in undefined behavior.</p> <p>In cases where default condition variable attributes are appropriate, the macro <code>PTHREAD_COND_INITIALIZER</code> can be used to initialize condition variables that are statically allocated. The effect is equivalent to dynamic initialization by a call to <code>pthread_cond_init()</code> with parameter <i>attr</i> specified as NULL, except that no error checks are performed.</p>				
RETURN VALUES	If successful, the <code>pthread_cond_init()</code> and <code>pthread_cond_destroy()</code> functions return 0. Otherwise, an error number is returned to indicate the error. The <code>EBUSY</code> and <code>EINVAL</code> error checks, if implemented, act as if they were performed immediately at the beginning of processing for the function and caused an error return prior to modifying the state of the condition variable specified by <i>cond</i> .				
ERRORS	<p>The <code>pthread_cond_init()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>EAGAIN</code></td> <td>The system lacked the necessary resources (other than memory) to initialize another condition variable.</td> </tr> <tr> <td><code>ENOMEM</code></td> <td>Insufficient memory exists to initialize the condition variable.</td> </tr> </table> <p>The <code>pthread_cond_init()</code> function may fail if:</p>	<code>EAGAIN</code>	The system lacked the necessary resources (other than memory) to initialize another condition variable.	<code>ENOMEM</code>	Insufficient memory exists to initialize the condition variable.
<code>EAGAIN</code>	The system lacked the necessary resources (other than memory) to initialize another condition variable.				
<code>ENOMEM</code>	Insufficient memory exists to initialize the condition variable.				

pthread_cond_init(3C)

EBUSY The implementation has detected an attempt to re-initialize the object referenced by *cond*, a previously initialized, but not yet destroyed, condition variable.

EINVAL The value specified by *attr* is invalid.

The `pthread_cond_destroy()` function may fail if:

EBUSY The implementation has detected an attempt to destroy the object referenced by *cond* while it is referenced (for example, while being used in a `pthread_cond_wait()` or `pthread_cond_timedwait()`) by another thread.

EINVAL The value specified by *cond* is invalid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `pthread_cond_signal(3C)`, `pthread_cond_broadcast(3C)`, `pthread_cond_wait(3C)`, `pthread_cond_timedwait(3C)`, `pthread_condattr_init(3C)`, `attributes(5)`, `condition(5)`, `standards(5)`

NAME	pthread_cond_signal, pthread_cond_broadcast – signal or broadcast a condition
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_cond_signal(pthread_cond_t *<i>cond</i>) ; int pthread_cond_broadcast(pthread_cond_t *<i>cond</i>) ;</pre>
DESCRIPTION	<p>These two functions are used to unblock threads blocked on a condition variable.</p> <p>The <code>pthread_cond_signal()</code> call unblocks at least one of the threads that are blocked on the specified condition variable <i>cond</i> (if any threads are blocked on <i>cond</i>).</p> <p>The <code>pthread_cond_broadcast()</code> call unblocks all threads currently blocked on the specified condition variable <i>cond</i>.</p> <p>If more than one thread is blocked on a condition variable, the scheduling policy determines the order in which threads are unblocked. When each thread unblocked as a result of a <code>pthread_cond_signal()</code> or <code>pthread_cond_broadcast()</code> returns from its call to <code>pthread_cond_wait()</code> or <code>pthread_cond_timedwait()</code>, the thread owns the mutex with which it called <code>pthread_cond_wait()</code> or <code>pthread_cond_timedwait()</code>. The thread(s) that are unblocked contend for the mutex according to the scheduling policy (if applicable), and as if each had called <code>pthread_mutex_lock()</code>.</p> <p>The <code>pthread_cond_signal()</code> or <code>pthread_cond_broadcast()</code> functions may be called by a thread whether or not it currently owns the mutex that threads calling <code>pthread_cond_wait()</code> or <code>pthread_cond_timedwait()</code> have associated with the condition variable during their waits; however, if predictable scheduling behavior is required, then that mutex is locked by the thread calling <code>pthread_cond_signal()</code> or <code>pthread_cond_broadcast()</code>.</p> <p>The <code>pthread_cond_signal()</code> and <code>pthread_cond_broadcast()</code> functions have no effect if there are no threads currently blocked on <i>cond</i>.</p>
RETURN VALUES	If successful, the <code>pthread_cond_signal()</code> and <code>pthread_cond_broadcast()</code> functions return 0. Otherwise, an error number is returned to indicate the error.
ERRORS	<p>The <code>pthread_cond_signal()</code> and <code>pthread_cond_broadcast()</code> function may fail if:</p> <p>EINVAL The value <i>cond</i> does not refer to an initialized condition variable.</p>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

pthread_cond_signal(3C)

SEE ALSO | [pthread_cond_init\(3C\)](#), [pthread_cond_wait\(3C\)](#),
[pthread_cond_timedwait\(3C\)](#), [attributes\(5\)](#), [condition\(5\)](#), [standards\(5\)](#)

NAME	pthread_cond_wait, pthread_cond_timedwait, pthread_cond_reltimedwait_np – wait on a condition
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_cond_wait(pthread_cond_t *restrict <i>cond</i>, pthread_mutex_t *restrict <i>mutex</i>); int pthread_cond_timedwait(pthread_cond_t *restrict <i>cond</i>, pthread_mutex_t *restrict <i>mutex</i>, const struct timespec *restrict <i>abstime</i>); int pthread_cond_reltimedwait_np(pthread_cond_t *<i>cond</i>, pthread_mutex_t *<i>mutex</i>, const struct timespec *<i>reltime</i>);</pre>
DESCRIPTION	<p>The pthread_cond_wait(), pthread_cond_timedwait(), and pthread_cond_reltimedwait_np() functions are used to block on a condition variable. They are called with <i>mutex</i> locked by the calling thread or undefined behaviour will result.</p> <p>These functions atomically release <i>mutex</i> and cause the calling thread to block on the condition variable <i>cond</i>. Atomically here means “atomically with respect to access by another thread to the mutex and then the condition variable”. That is, if another thread is able to acquire the mutex after the about-to-block thread has released it, then a subsequent call to pthread_cond_signal() or pthread_cond_broadcast() in that thread behaves as if it were issued after the about-to-block thread has blocked.</p> <p>Upon successful return, the mutex has been locked and is owned by the calling thread.</p> <p>When using condition variables there is always a boolean predicate, an invariant, associated with each condition wait that must be true before the thread should proceed. Spurious wakeups from the pthread_cond_wait(), pthread_cond_timedwait(), or pthread_cond_reltimedwait_np() functions may occur. Since the return from pthread_cond_wait(), pthread_cond_timedwait(), or pthread_cond_reltimedwait_np() does not imply anything about the value of this predicate, the predicate should always be re-evaluated.</p> <p>The order in which blocked threads are awakened by pthread_cond_signal() or pthread_cond_broadcast() is determined by the scheduling policy. See pthreads(5).</p> <p>The effect of using more than one mutex for concurrent pthread_cond_wait(), pthread_cond_timedwait(), or pthread_cond_reltimedwait_np() operations on the same condition variable will result in undefined behavior.</p> <p>A condition wait (whether timed or not) is a cancellation point. When the cancelability enable state of a thread is set to PTHREAD_CANCEL_DEFERRED, a side effect of acting upon a cancellation request while in a condition wait is that the mutex is re-acquired before calling the first cancellation cleanup handler.</p>

pthread_cond_wait(3C)

A thread that has been unblocked because it has been canceled while blocked in a call to `pthread_cond_wait()` or `pthread_cond_timedwait()` does not consume any condition signal that may be directed concurrently at the condition variable if there are other threads blocked on the condition variable.

The `pthread_cond_timedwait()` function is the same as `pthread_cond_wait()` except that an error is returned if the absolute time specified by *abstime* passes (that is, system time equals or exceeds *abstime*) before the condition *cond* is signaled or broadcasted, or if the absolute time specified by *abstime* has already been passed at the time of the call. When such time-outs occur, `pthread_cond_timedwait()` will nonetheless release and reacquire the mutex referenced by *mutex*. The function `pthread_cond_timedwait()` is also a cancellation point.

The `pthread_cond_reltimedwait_np()` function is a non-standard extension provided by the Solaris version of pthreads as indicated by the “_np” (non-portable) suffix. The `pthread_cond_reltimedwait_np()` function is the same as `pthread_cond_timedwait()` except that the *reltime* argument specifies a non-negative time relative to the current system time rather than an absolute time. An error value is returned if the relative time passes (that is, system time equals or exceeds the starting system time plus the relative time) before the condition *cond* is signaled or broadcasted. When such timeouts occur, `pthread_cond_reltimedwait_np()` releases and reacquires the mutex referenced by *mutex*. The `pthread_cond_reltimedwait_np()` function is also a cancellation point.

If a signal is delivered to a thread waiting for a condition variable, upon return from the signal handler the thread resumes waiting for the condition variable as if it was not interrupted, or it returns 0 due to spurious wakeup.

RETURN VALUES

Except in the case of `ETIMEDOUT`, all these error checks act as if they were performed immediately at the beginning of processing for the function and cause an error return, in effect, prior to modifying the state of the mutex specified by *mutex* or the condition variable specified by *cond*.

Upon successful completion, 0 is returned. Otherwise, an error value is returned to indicate the error.

ERRORS

The `pthread_cond_timedwait()` function will fail if:

`ETIMEDOUT` The absolute time specified by *abstime* to `pthread_cond_timedwait()` has passed.

The `pthread_cond_reltimedwait_np()` function will fail if:

`EINVAL` The value specified by *reltime* is invalid.

`ETIMEDOUT` The relative time specified by *reltime* to `pthread_cond_reltimedwait_np()` has passed.

The `pthread_cond_wait()` and `pthread_cond_timedwait()` functions may fail if:

EINVAL	The value specified by <i>cond</i> , <i>mutex</i> , or <i>abstime</i> is invalid.
EINVAL	Different mutexes were supplied for concurrent <code>pthread_cond_wait()</code> or <code>pthread_cond_timedwait()</code> , operations on the same condition variable.
EPERM	The mutex was not owned by the current thread at the time of the call.

When a thread makes a call to `pthread_cond_wait()`, `pthread_cond_timedwait()`, or `pthread_cond_reltimedwait_np()`, if the symbol `_POSIX_THREAD_PRIO_INHERIT` is defined and the mutex is initialized with the protocol attribute having the value `PTHREAD_PRIO_INHERIT` and the robustness attribute having the value `PTHREAD_MUTEX_ROBUST_NP` (see [pthread_mutexattr_getrobust_np\(3C\)](#)), the `pthread_cond_wait()`, `pthread_cond_timedwait()`, and `pthread_cond_reltimedwait_np()` functions will fail if:

EOWNERDEAD	The last owner of this mutex died while holding the mutex. This mutex is now owned by the caller. The caller must now attempt to make the state protected by the mutex consistent. If it is able to clean up the state, then it should call <code>pthread_mutex_consistent_np()</code> for the mutex and unlock the mutex. Subsequent calls to <code>pthread_cond_wait()</code> , <code>pthread_cond_timedwait()</code> , and <code>pthread_cond_reltimedwait_np()</code> will behave normally, as before. If the caller is not able to clean up the state, <code>pthread_mutex_consistent_np()</code> should not be called for the mutex, but it should be unlocked. Subsequent calls to <code>pthread_cond_wait()</code> , <code>pthread_cond_timedwait()</code> , and <code>pthread_cond_reltimedwait_np()</code> will fail to acquire the mutex with the error value <code>ENOTRECOVERABLE</code> . If the owner who acquired the lock with <code>EOWNERDEAD</code> dies, the next owner will acquire the lock with <code>EOWNERDEAD</code> .
ENOTRECOVERABLE	The mutex trying to be acquired is protecting the state that has been left irrecoverable by the mutex's last owner, who died while holding the lock. The mutex has not been acquired. This condition can occur when the lock was previously acquired with <code>EOWNERDEAD</code> , and the owner was not able to clean up the state and unlocked the mutex without making the mutex consistent.

pthread_cond_wait(3C)

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [pthread_cond_signal\(3C\)](#), [pthread_cond_broadcast\(3C\)](#), [attributes\(5\)](#), [condition\(5\)](#), [pthreads\(5\)](#), [standards\(5\)](#)

NAME	pthread_create – create a thread
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_create(pthread_t *restrict <i>thread</i>, const pthread_attr_t *restrict <i>attr</i>, void *(*<i>start_routine</i>)(void*), void *restrict <i>arg</i>);</pre>
DESCRIPTION	<p>The <code>pthread_create()</code> function is used to create a new thread, with attributes specified by <i>attr</i>, within a process. If <i>attr</i> is <code>NULL</code>, the default attributes are used. (See pthread_attr_init(3C)). If the attributes specified by <i>attr</i> are modified later, the thread's attributes are not affected. Upon successful completion, <code>pthread_create()</code> stores the ID of the created thread in the location referenced by <i>thread</i>.</p> <p>The thread is created executing <i>start_routine</i> with <i>arg</i> as its sole argument. If the <i>start_routine</i> returns, the effect is as if there was an implicit call to <code>pthread_exit()</code> using the return value of <i>start_routine</i> as the exit status. Note that the thread in which <code>main()</code> was originally invoked differs from this. When it returns from <code>main()</code>, the effect is as if there was an implicit call to <code>exit()</code> using the return value of <code>main()</code> as the exit status.</p> <p>The signal state of the new thread is initialised as follows:</p> <ul style="list-style-type: none"> ■ The signal mask is inherited from the creating thread. ■ The set of signals pending for the new thread is empty. <p>Default thread creation:</p> <pre>pthread_t tid; void *start_func(void *), *arg; pthread_create(&tid, NULL, start_func, arg);</pre> <p>This would have the same effect as:</p> <pre>pthread_attr_t attr; pthread_attr_init(&attr); /* initialize attr with default attributes */ pthread_create(&tid, &attr, start_func, arg);</pre> <p>User-defined thread creation: To create a thread that is scheduled on a system-wide basis, use:</p> <pre>pthread_attr_init(&attr); /* initialize attr with default attributes */ pthread_attr_setscope(&attr, PTHREAD_SCOPE_SYSTEM); /* system-wide contention */ pthread_create(&tid, &attr, start_func, arg);</pre> <p>To customize the attributes for POSIX threads, see pthread_attr_init(3C).</p>

pthread_create(3C)

A new thread created with `pthread_create()` uses the stack specified by the `stackaddr` attribute, and the stack continues for the number of bytes specified by the `stacksize` attribute. By default, the stack size is 1 megabyte for 32-bit processes and 2 megabyte for 64-bit processes (see `pthread_attr_setstacksize(3C)`). If the default is used for both the `stackaddr` and `stacksize` attributes, `pthread_create()` creates a stack for the new thread with at least 1 megabyte for 32-bit processes and 2 megabyte for 64-bit processes. (For customizing stack sizes, see NOTES).

If `pthread_create()` fails, no new thread is created and the contents of the location referenced by `thread` are undefined.

RETURN VALUES If successful, the `pthread_create()` function returns 0. Otherwise, an error number is returned to indicate the error.

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

EAGAIN The system lacked the necessary resources to create another thread, or the system-imposed limit on the total number of threads in a process `PTHREAD_THREADS_MAX` would be exceeded.

EINVAL The value specified by `attr` is invalid.

EPERM The caller does not have appropriate permission to set the required scheduling parameters or scheduling policy.

EXAMPLES **EXAMPLE 1** Example of concurrency with multithreading

The following is an example of concurrency with multithreading. Since POSIX threads and Solaris threads are fully compatible even within the same process, this example uses `pthread_create()` if you execute `a.out 0`, or `thr_create()` if you execute `a.out 1`.

Five threads are created that simultaneously perform a time-consuming function, `sleep(10)`. If the execution of this process is timed, the results will show that all five individual calls to `sleep` for ten-seconds completed in about ten seconds, even on a uniprocessor. If a single-threaded process calls `sleep(10)` five times, the execution time will be about 50-seconds.

The command-line to time this process is:

POSIX threading `/usr/bin/time a.out 0`

Solaris threading `/usr/bin/time a.out 1`

```
/* cc thisfile.c -lthread -lpthread */
#define _REENTRANT /* basic 3-lines for threads */
#include <pthread.h>
#include <thread.h>

#define NUM_THREADS 5
#define SLEEP_TIME 10

void *sleeping(void *); /* thread routine */
```

EXAMPLE 1 Example of concurrency with multithreading (Continued)

```

int i;
thread_t tid[NUM_THREADS];      /* array of thread IDs */

int
main(int argc, char *argv[])
{
    if (argc == 1) {
        printf("use 0 as arg1 to use pthread_create( )\n");
        printf("or use 1 as arg1 to use thr_create( )\n");
        return (1);
    }

    switch (*argv[1]) {
case '0': /* POSIX */
        for ( i = 0; i < NUM_THREADS; i++)
            pthread_create(&tid[i], NULL, sleeping,
                (void *)SLEEP_TIME);
        for ( i = 0; i < NUM_THREADS; i++)
            pthread_join(tid[i], NULL);
        break;

case '1': /* Solaris */
        for ( i = 0; i < NUM_THREADS; i++)
            thr_create(NULL, 0, sleeping, (void *)SLEEP_TIME, 0,
                &tid[i]);
        while (thr_join(0, NULL, NULL) == 0)
            ;
        break;
    } /* switch */
    printf("main( ) reporting that all %d threads have terminated\n", i);
    return (0);
} /* main */

void *
sleeping(void *arg)
{
    int sleep_time = (int)arg;
    printf("thread %d sleeping %d seconds ...\n", thr_self( ), sleep_time);
    sleep(sleep_time);
    printf("\nthread %d awakening\n", thr_self( ));
    return (NULL);
}

```

If `main()` had not waited for the completion of the other threads (using [pthread_join\(3C\)](#) or [thr_join\(3C\)](#)), it would have continued to process concurrently until it reached the end of its routine and the entire process would have exited prematurely. See [exit\(2\)](#).

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

pthread_create(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [fork\(2\)](#), [pthread_attr_init\(3C\)](#), [pthread_cancel\(3C\)](#), [pthread_exit\(3C\)](#), [pthread_join\(3C\)](#), [sysconf\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NOTES Multithreaded application threads execute independently of each other, so their relative behavior is unpredictable. Therefore, it is possible for the thread executing `main()` to finish before all other user application threads. The [pthread_join\(3C\)](#) function, on the other hand, must specify the terminating thread (IDs) for which it will wait.

A user-specified stack size must be greater than the value `PTHREAD_STACK_MIN`. A minimum stack size may not accommodate the stack frame for the user thread function `start_func`. If a stack size is specified, it must accommodate `start_func` requirements and the functions that it may call in turn, in addition to the minimum requirement.

It is usually very difficult to determine the runtime stack requirements for a thread. `PTHREAD_STACK_MIN` specifies how much stack storage is required to execute a `NULL start_func`. The total runtime requirements for stack storage are dependent on the storage required to do runtime linking, the amount of storage required by library runtimes (as `printf()`) that your thread calls. Since these storage parameters are not known before the program runs, it is best to use default stacks. If you know your runtime requirements or decide to use stacks that are larger than the default, then it makes sense to specify your own stacks.

NAME	pthread_detach – detach a thread						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_detach(pthread_t <i>thread</i>);</pre>						
DESCRIPTION	<p>The <code>pthread_detach()</code> function is used to indicate to the implementation that storage for the thread <i>thread</i> can be reclaimed when that thread terminates. In other words, <code>pthread_detach()</code> dynamically resets the <i>detachstate</i> attribute of the thread to <code>PTHREAD_CREATE_DETACHED</code>. After a successful call to this function, it would not be necessary to reclaim the thread using <code>pthread_join()</code>. See pthread_join(3C). If <i>thread</i> has not terminated, <code>pthread_detach()</code> will not cause it to terminate. The effect of multiple <code>pthread_detach()</code> calls on the same target thread is unspecified.</p>						
RETURN VALUES	If successful, <code>pthread_detach()</code> returns 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	<p>The <code>pthread_detach()</code> function will fail if:</p> <table> <tr> <td>EINVAL</td> <td>The implementation has detected that the value specified by <i>thread</i> does not refer to a joinable thread.</td> </tr> <tr> <td>ESRCH</td> <td>No thread could be found corresponding to that specified by the given thread ID.</td> </tr> </table>	EINVAL	The implementation has detected that the value specified by <i>thread</i> does not refer to a joinable thread.	ESRCH	No thread could be found corresponding to that specified by the given thread ID.		
EINVAL	The implementation has detected that the value specified by <i>thread</i> does not refer to a joinable thread.						
ESRCH	No thread could be found corresponding to that specified by the given thread ID.						
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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	pthread_create(3C) , pthread_join(3C) , attributes(5) , standards(5)						

pthread_equal(3C)

NAME pthread_equal – compare thread IDs

SYNOPSIS

```
cc -mt [ flag... ] file... -pthread [ -lrt library... ]
#include <pthread.h>

int pthread_equal(pthread_t t1, pthread_t t2);
```

DESCRIPTION The pthread_equal() function compares the thread IDs *t1* and *t2*.

RETURN VALUES The pthread_equal() function returns a non-zero value if *t1* and *t2* are equal. Otherwise, 0 is returned.

If *t1* or *t2* is an invalid thread ID, the behavior is undefined.

ERRORS No errors are defined.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO pthread_create(3C), pthread_self(3C), attributes(5), standards(5)

NOTES Solaris thread IDs do not require an equivalent function because the thread_t structure is an unsigned int.

NAME	pthread_exit – terminate calling thread						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -pthread [-lrt <i>library...</i>] #include <pthread.h> void pthread_exit(void *<i>value_ptr</i>);</pre>						
DESCRIPTION	<p>The pthread_exit() function terminates the calling thread, in a similar way that exit(3C) terminates the calling process. If the thread is not detached, the exit status specified by <i>value_ptr</i> is made available to any successful join with the terminating thread. See pthread_join(3C). Any cancellation cleanup handlers that have been pushed and not yet popped are popped in the reverse order that they were pushed and then executed. After all cancellation cleanup handlers have been executed, if the thread has any thread-specific data, appropriate destructor functions will be called in an unspecified order. Thread termination does not release any application visible process resources, including, but not limited to, mutexes and file descriptors, nor does it perform any process level cleanup actions, including, but not limited to, calling any atexit() routines that might exist.</p> <p>An implicit call to pthread_exit() is made when a thread other than the thread in which main() was first invoked returns from the start routine that was used to create it. The function's return value serves as the thread's exit status.</p> <p>The behavior of pthread_exit() is undefined if called from a cancellation cleanup handler or destructor function that was invoked as a result of either an implicit or explicit call to pthread_exit().</p> <p>After a thread has terminated, the result of access to local (auto) variables of the thread is undefined. Thus, references to local variables of the exiting thread should not be used for the pthread_exit() <i>value_ptr</i> parameter value.</p> <p>The process exits with an exit status of 0 after the last thread has been terminated. The behavior is as if the implementation called exit() with a 0 argument at thread termination time.</p>						
RETURN VALUES	The pthread_exit() function cannot return to its caller.						
ERRORS	No errors are defined.						
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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	exit(3C), pthread_cancel(3C), pthread_create(3C), pthread_join(3C), pthread_key_create(3C), attributes(5), standards(5)						

pthread_getconcurrency(3C)

NAME	pthread_getconcurrency, pthread_setconcurrency – get or set level of concurrency				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_getconcurrency(void); int pthread_setconcurrency(int <i>new_level</i>);</pre>				
DESCRIPTION	<p>Unbound threads in a process may or may not be required to be simultaneously active. By default, the threads implementation ensures that a sufficient number of threads are active so that the process can continue to make progress. While this conserves system resources, it may not produce the most effective level of concurrency.</p> <p>The <code>pthread_setconcurrency()</code> function allows an application to inform the threads implementation of its desired concurrency level, <i>new_level</i>. The actual level of concurrency provided by the implementation as a result of this function call is unspecified.</p> <p>If <i>new_level</i> is 0, it causes the implementation to maintain the concurrency level at its discretion as if <code>pthread_setconcurrency()</code> was never called.</p> <p>The <code>pthread_getconcurrency()</code> function returns the value set by a previous call to the <code>pthread_setconcurrency()</code> function. If the <code>pthread_setconcurrency()</code> function was not previously called, this function returns 0 to indicate that the implementation is maintaining the concurrency level.</p> <p>When an application calls <code>pthread_setconcurrency()</code> it is informing the implementation of its desired concurrency level. The implementation uses this as a hint, not a requirement.</p> <p>If an implementation does not support multiplexing of user threads on top of several kernel scheduled entities, the <code>pthread_setconcurrency()</code> and <code>pthread_getconcurrency()</code> functions will be provided for source code compatibility but they will have no effect when called. To maintain the function semantics, the <i>new_level</i> parameter will be saved when <code>pthread_setconcurrency()</code> is called so that a subsequent call to <code>pthread_getconcurrency()</code> returns the same value.</p>				
RETURN VALUES	<p>If successful, the <code>pthread_setconcurrency()</code> function returns 0. Otherwise, an error number is returned to indicate the error.</p> <p>The <code>pthread_getconcurrency()</code> function always returns the concurrency level set by a previous call to <code>pthread_setconcurrency()</code>. If the <code>pthread_setconcurrency()</code> function has never been called, <code>pthread_getconcurrency()</code> returns 0.</p>				
ERRORS	<p>The <code>pthread_setconcurrency()</code> function will fail if:</p> <table><tr><td>EINVAL</td><td>The value specified by <i>new_level</i> is negative.</td></tr><tr><td>EAGAIN</td><td>The value specific by <i>new_level</i> would cause a system resource to be exceeded.</td></tr></table>	EINVAL	The value specified by <i>new_level</i> is negative.	EAGAIN	The value specific by <i>new_level</i> would cause a system resource to be exceeded.
EINVAL	The value specified by <i>new_level</i> is negative.				
EAGAIN	The value specific by <i>new_level</i> would cause a system resource to be exceeded.				

`pthread_getconcurrency(3C)`

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `pthread_create(3C)`, `pthread_attr_init(3C)`, `attributes(5)`, `standards(5)`

pthread_getschedparam(3C)

NAME	pthread_getschedparam, pthread_setschedparam – access dynamic thread scheduling parameters
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_getschedparam(pthread_t <i>thread</i>, int *restrict <i>policy</i>, struct sched_param *restrict <i>param</i>); int pthread_setschedparam(pthread_t <i>thread</i>, int <i>policy</i>, const struct sched_param *<i>param</i>);</pre>
DESCRIPTION	<p>The pthread_getschedparam() and pthread_setschedparam() functions allow the scheduling policy and scheduling parameters of individual threads within a multithreaded process to be retrieved and set. Supported policies are SCHED_FIFO, SCHED_RR, and SCHED_OTHER. See pthreads(5). For SCHED_FIFO, SCHED_RR, and SCHED_OTHER, the affected scheduling parameter is the sched_priority member of the sched_param structure.</p> <p>The pthread_getschedparam() function retrieves the scheduling policy and scheduling parameters for the thread whose thread ID is given by <i>thread</i> and stores those values in <i>policy</i> and <i>param</i>, respectively. The priority value returned from pthread_getschedparam() is the value specified by the most recent pthread_setschedparam() or pthread_create() call affecting the target thread, and reflects any temporary adjustments to its priority as a result of any priority inheritance or ceiling functions. The pthread_setschedparam() function sets the scheduling policy and associated scheduling parameters for the thread whose thread ID is given by <i>thread</i> to the policy and associated parameters provided in <i>policy</i> and <i>param</i>, respectively.</p> <p>If the pthread_setschedparam() function fails, no scheduling parameters will be changed for the target thread.</p>
RETURN VALUES	If successful, the pthread_getschedparam() and pthread_setschedparam() functions return 0. Otherwise, an error number is returned to indicate the error.
ERRORS	The pthread_getschedparam() function may fail if: ESRCH The value specified by <i>thread</i> does not refer to a existing thread.
	The pthread_setschedparam() function may fail if: EINVAL The value specified by <i>policy</i> or one of the scheduling parameters associated with the scheduling policy <i>policy</i> is invalid. EPERM The caller does not have the appropriate permission to set either the scheduling parameters or the scheduling policy of the specified thread. ESRCH The value specified by <i>thread</i> does not refer to a existing thread.

pthread_getschedparam(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `pthread_attr_init(3C)`, `sched_getparam(3RT)`, `sched_setparam(3RT)`, `sched_getscheduler(3RT)`, `sched_setscheduler(3RT)`, `attributes(5)`, `pthreads(5)`, `standards(5)`

pthread_getspecific(3C)

NAME	pthread_getspecific, pthread_setspecific – manage thread-specific data						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_setspecific(pthread_key_t <i>key</i>, const void *<i>value</i>); void *pthread_getspecific(pthread_key_t <i>key</i>);</pre>						
DESCRIPTION	<p>The <code>pthread_setspecific()</code> function associates a thread-specific <i>value</i> with a <i>key</i> obtained by way of a previous call to <code>pthread_key_create()</code>. Different threads may bind different values to the same key. These values are typically pointers to blocks of dynamically allocated memory that have been reserved for use by the calling thread.</p> <p>The <code>pthread_getspecific()</code> function returns the value currently bound to the specified <i>key</i> on behalf of the calling thread.</p> <p>The effect of calling <code>pthread_setspecific()</code> or <code>pthread_getspecific()</code> with a <i>key</i> value not obtained from <code>pthread_key_create()</code> or after <i>key</i> has been deleted with <code>pthread_key_delete()</code> is undefined.</p> <p>Both <code>pthread_setspecific()</code> and <code>pthread_getspecific()</code> may be called from a thread-specific data destructor function. However, calling <code>pthread_setspecific()</code> from a destructor may result in lost storage or infinite loops.</p>						
RETURN VALUES	<p>The <code>pthread_getspecific()</code> function returns the thread-specific data value associated with the given <i>key</i>. If no thread-specific data value is associated with <i>key</i>, then the value NULL is returned.</p> <p>Upon successful completion, the <code>pthread_setspecific()</code> function returns 0. Otherwise, an error number is returned to indicate the error.</p>						
ERRORS	<p>The <code>pthread_setspecific()</code> function will fail if:</p> <table><tr><td>ENOMEM</td><td>Insufficient memory exists to associate the value with the key.</td></tr></table> <p>The <code>pthread_setspecific()</code> function may fail if:</p> <table><tr><td>EINVAL</td><td>The key value is invalid.</td></tr></table> <p>The <code>pthread_getspecific()</code> function does not return errors.</p>	ENOMEM	Insufficient memory exists to associate the value with the key.	EINVAL	The key value is invalid.		
ENOMEM	Insufficient memory exists to associate the value with the key.						
EINVAL	The key value is invalid.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>pthread_key_create(3C)</code> <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	pthread_join – wait for thread termination						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_join(pthread_t <i>thread</i>, void **<i>status</i>);</pre>						
DESCRIPTION	<p>The pthread_join() function suspends processing of the calling thread until the target <i>thread</i> completes. <i>thread</i> must be a member of the current process and it cannot be a detached thread. See pthread_create(3C).</p> <p>If two or more threads wait for the same thread to complete, all will suspend processing until the thread has terminated, and then one thread will return successfully and the others will return with an error of ESRCH. The pthread_join() function will not block processing of the calling thread if the target <i>thread</i> has already terminated.</p> <p>If a pthread_join() call returns successfully with a non-null <i>status</i> argument, the value passed to pthread_exit(3C) by the terminating thread will be placed in the location referenced by <i>status</i>.</p> <p>If the pthread_join() calling thread is cancelled, then the target <i>thread</i> will remain joinable by pthread_join(). However, the calling thread may set up a cancellation cleanup handler on <i>thread</i> prior to the join call, which may detach the target <i>thread</i> by calling pthread_detach(3C). See pthread_detach(3C) and pthread_cancel(3C).</p>						
RETURN VALUES	If successful, pthread_join() returns 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	<table border="0"> <tr> <td style="vertical-align: top;">EDEADLK</td> <td>A joining deadlock would occur, such as when a thread attempts to wait for itself.</td> </tr> <tr> <td style="vertical-align: top;">EINVAL</td> <td>The thread corresponding to the given thread ID is a detached thread.</td> </tr> <tr> <td style="vertical-align: top;">ESRCH</td> <td>No thread could be found corresponding to the given thread ID.</td> </tr> </table>	EDEADLK	A joining deadlock would occur, such as when a thread attempts to wait for itself.	EINVAL	The thread corresponding to the given thread ID is a detached thread.	ESRCH	No thread could be found corresponding to the given thread ID.
EDEADLK	A joining deadlock would occur, such as when a thread attempts to wait for itself.						
EINVAL	The thread corresponding to the given thread ID is a detached thread.						
ESRCH	No thread could be found corresponding to the given thread ID.						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	pthread_cancel(3C), pthread_create(3C), pthread_detach(3C), pthread_exit(3C), wait(3C), attributes(5), standards(5)						
NOTES	The pthread_join(3C) function must specify the <i>thread</i> ID for whose termination it will wait.						

pthread_join(3C)

Calling `pthread_join()` also "detaches" the thread; that is, `pthread_join()` includes the effect of the `pthread_detach()` function. If a thread were to be cancelled when blocked in `pthread_join()`, an explicit detach would have to be performed in the cancellation cleanup handler. The `pthread_detach()` function exists primarily for this purpose.

NAME	pthread_key_create – create thread-specific data key
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -pthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_key_create(pthread_key_t *<i>key</i>, void (*<i>destructor</i>, void*));</pre>
DESCRIPTION	<p>This function creates a thread-specific data key visible to all threads in the process. Key values provided by <code>pthread_key_create()</code> are opaque objects used to locate thread-specific data. Although the same key value may be used by different threads, the values bound to the key by <code>pthread_setspecific()</code> are maintained on a per-thread basis and persist for the life of the calling thread.</p> <p>Upon key creation, the value <code>NULL</code> is associated with the new key in all active threads. Upon thread creation, the value <code>NULL</code> is associated with all defined keys in the new thread.</p> <p>An optional destructor function may be associated with each key value. At thread exit, if a key value has a non-<code>NULL</code> destructor pointer, and the thread has a non-<code>NULL</code> value associated with that key, the function pointed to is called with the current associated value as its sole argument. Destructors can be called in any order.</p> <p>If, after all the destructors have been called for all keys with non-<code>NULL</code> values, there are still some keys with non-<code>NULL</code> values, the process will be repeated. If, after at least <code>PTHREAD_DESTRUCTOR_ITERATIONS</code> iterations of destructor calls for outstanding non-<code>NULL</code> values, there are still some keys with non-<code>NULL</code> values, the process is continued, even though this might result in an infinite loop.</p>
RETURN VALUES	If successful, the <code>pthread_key_create()</code> function stores the newly created key value at <i>key</i> and returns 0. Otherwise, an error number is returned to indicate the error.
ERRORS	<p>The <code>pthread_key_create()</code> function will fail if:</p> <p>EAGAIN The system lacked the necessary resources to create another thread-specific data key, or the system-imposed limit on the total number of keys per process <code>PTHREAD_KEYS_MAX</code> has been exceeded.</p> <p>ENOMEM Insufficient memory exists to create the key.</p> <p>The <code>pthread_key_create()</code> function will not return an error code of <code>EINTR</code>.</p>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

pthread_key_create(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [pthread_getspecific\(3C\)](#), [pthread_setspecific\(3C\)](#),
[pthread_key_delete\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	pthread_key_delete – delete thread-specific data key						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -pthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_key_delete(pthread_key_t <i>key</i>);</pre>						
DESCRIPTION	<p>The pthread_key_delete() function deletes a thread-specific data key previously returned by pthread_key_create(). The thread-specific data values associated with <i>key</i> need not be NULL at the time pthread_key_delete() is called. It is the responsibility of the application to free any application storage or perform any cleanup actions for data structures related to the deleted key or associated thread-specific data in any threads; this cleanup can be done either before or after pthread_key_delete() is called. Any attempt to use <i>key</i> following the call to pthread_key_delete() results in undefined behaviour.</p> <p>The pthread_key_delete() function is callable from within destructor functions. No destructor functions will be invoked by pthread_key_delete(). Any destructor function that may have been associated with <i>key</i> will no longer be called upon thread exit.</p>						
RETURN VALUES	If successful, the pthread_key_delete() function returns 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	<p>The pthread_key_delete() function may fail if:</p> <p>EINVAL The <i>key</i> value is invalid.</p> <p>The pthread_key_delete() function will not return an error code of EINTR.</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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	pthread_key_create(3C), attributes(5), standards(5)						

pthread_kill(3C)

NAME pthread_kill – send a signal to a thread

SYNOPSIS

```
cc -mt [ flag... ] file... -pthread [ -lrt library... ]
#include <signal.h>
#include <pthread.h>

int pthread_kill(pthread_t thread, int sig);
```

DESCRIPTION The `pthread_kill()` function is used to request that a signal be delivered to the specified thread.

As in `kill()`, if *sig* is 0, error checking is performed but no signal is actually sent.

RETURN VALUES Upon successful completion, the function returns a value of 0. Otherwise the function returns an error number. If the `pthread_kill()` function fails, no signal is sent.

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

ESRCH	No thread could be found corresponding to that specified by the given thread ID.
EINVAL	The value of the <i>sig</i> argument is an invalid or unsupported signal number.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `kill(1)`, `pthread_self(3C)`, `pthread_sigmask(3C)`, `raise(3C)`, `attributes(5)`, `standards(5)`

pthread_mutexattr_getprioceiling(3C)

NAME	pthread_mutexattr_getprioceiling, pthread_mutexattr_setprioceiling – get and set prioceiling attribute of mutex attribute object
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutexattr_getprioceiling(const pthread_mutexattr_t *restrict <i>attr</i>, int *restrict <i>prioceiling</i>); int pthread_mutexattr_setprioceiling(pthread_mutexattr_t *<i>attr</i>, int <i>prioceiling</i> int *<i>oldceiling</i>);</pre>
DESCRIPTION	<p>The pthread_mutexattr_getprioceiling() and pthread_mutexattr_setprioceiling() functions, respectively, get and set the priority ceiling attribute of a mutex attribute object pointed to by <i>attr</i>, which was previously created by the pthread_mutexattr_init() function.</p> <p>The <i>prioceiling</i> attribute contains the priority ceiling of initialized mutexes. The values of <i>prioceiling</i> must be within the maximum range of priorities defined by SCHED_FIFO.</p> <p>The <i>prioceiling</i> attribute defines the priority ceiling of initialized mutexes, which is the minimum priority level at which the critical section guarded by the mutex is executed. In order to avoid priority inversion, the priority ceiling of the mutex must be set to a priority higher than or equal to the highest priority of all the threads that may lock that mutex. The values of <i>prioceiling</i> must be within the maximum range of priorities defined under the SCHED_FIFO scheduling policy.</p> <p>The ceiling value should be drawn from the range of priorities for the SCHED_FIFO policy. When a thread acquires such a mutex, the policy of the thread at mutex acquisition should match that from which the ceiling value was derived (SCHED_FIFO, in this case). If a thread changes its scheduling policy while holding a ceiling mutex, the behavior of pthread_mutex_lock() and pthread_mutex_unlock() on this mutex is undefined. See pthread_mutex_lock(3C).</p> <p>The ceiling value should not be treated as a persistent value resident in a pthread_mutex_t that is valid across upgrades of Solaris. The semantics of the actual ceiling value are determined by the existing priority range for the SCHED_FIFO policy, as returned by the sched_get_priority_min(3RT) and sched_get_priority_max() functions when called on the version of Solaris on which the ceiling value is being utilized.</p>
RETURN VALUES	Upon successful completion, the pthread_mutexattr_getprioceiling() and pthread_mutexattr_setprioceiling() functions return 0. Otherwise, an error number is returned to indicate the error.
ERRORS	The pthread_mutexattr_getprioceiling() and pthread_mutexattr_setprioceiling() functions will fail if: ENOSYS The _POSIX_THREAD_PRIO_PROTECT option is not defined and the system does not support the function.

pthread_mutexattr_getprioceiling(3C)

The `pthread_mutexattr_getprioceiling()` function will fail if:

`EINVAL` The value specified by *attr* is `NULL`.

The `pthread_mutexattr_setprioceiling()` functions will fail if:

`EINVAL` The value specified by *attr* is `NULL` or *prioceiling* is invalid.

The `pthread_mutexattr_getprioceiling()` and `pthread_mutexattr_setprioceiling()` functions may fail if:

`EINVAL` The value specified by *attr* or *prioceiling* is invalid.

`EPERM` The caller does not have the privilege to perform the operation.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `pthread_cond_init(3C)`, `pthread_create(3C)`, `pthread_mutex_init(3C)`, `pthread_mutex_lock(3C)`, `sched_get_priority_min(3RT)`, `attributes(5)`, `standards(5)`

NAME	pthread_mutexattr_getprotocol, pthread_mutexattr_setprotocol – get and set protocol attribute of mutex attribute object
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutexattr_getprotocol(const pthread_mutexattr_t *restrict <i>attr</i>, int *restrict <i>protocol</i>); int pthread_mutexattr_setprotocol(pthread_mutexattr_t *<i>attr</i>, int <i>protocol</i>);</pre>
DESCRIPTION	<p>The <code>pthread_mutexattr_setprotocol()</code> and <code>pthread_mutexattr_getprotocol()</code> functions, respectively, set and get the protocol attribute of a mutex attribute object pointed to by <i>attr</i>, which was previously created by the <code>pthread_mutexattr_init()</code> function.</p> <p>The <i>protocol</i> attribute defines the protocol to be followed in utilizing mutexes. The value of <i>protocol</i> may be one of <code>PTHREAD_PRIO_NONE</code>, <code>PTHREAD_PRIO_INHERIT</code>, or <code>PTHREAD_PRIO_PROTECT</code>, which are defined by the header <code><pthread.h></code>.</p> <p>When a thread owns a mutex with the <code>PTHREAD_PRIO_NONE</code> protocol attribute, its priority and scheduling are not affected by its mutex ownership.</p> <p>When a thread is blocking higher priority threads because of owning one or more mutexes with the <code>PTHREAD_PRIO_INHERIT</code> protocol attribute, it executes at the higher of its priority or the priority of the highest priority thread waiting on any of the mutexes owned by this thread and initialized with this protocol.</p> <p>When a thread owns one or more mutexes initialized with the <code>PTHREAD_PRIO_PROTECT</code> protocol, it executes at the higher of its priority or the highest of the priority ceilings of all the mutexes owned by this thread and initialized with this attribute, regardless of whether other threads are blocked on any of these mutexes.</p> <p>While a thread is holding a mutex that has been initialized with the <code>PRIO_INHERIT</code> or <code>PRIO_PROTECT</code> protocol attributes, it will not be subject to being moved to the tail of the scheduling queue at its priority in the event that its original priority is changed, such as by a call to <code>sched_setparam()</code>. Likewise, when a thread unlocks a mutex that has been initialized with the <code>PRIO_INHERIT</code> or <code>PRIO_PROTECT</code> protocol attributes, it will not be subject to being moved to the tail of the scheduling queue at its priority in the event that its original priority is changed.</p> <p>If a thread simultaneously owns several mutexes initialized with different protocols, it will execute at the highest of the priorities that it would have obtained by each of these protocols.</p> <p>When a thread makes a call to <code>pthread_mutex_lock()</code>, if the symbol <code>_POSIX_THREAD_PRIO_INHERIT</code> is defined and the mutex was initialized with the protocol attribute having the value <code>PTHREAD_PRIO_INHERIT</code>, when the calling thread is blocked because the mutex is owned by another thread, that owner thread</p>

pthread_mutexattr_getprotocol(3C)

will inherit the priority level of the calling thread as long as it continues to own the mutex. The implementation updates its execution priority to the maximum of its assigned priority and all its inherited priorities. Furthermore, if this owner thread becomes blocked on another mutex, the same priority inheritance effect will be propagated to the other owner thread, in a recursive manner.

If the symbol `_POSIX_THREAD_PRIO_INHERIT` is defined, when a mutex initialized with the protocol attribute having the value `PTHREAD_PRIO_INHERIT` dies, the behavior depends on the robustness attribute of the mutex. See [pthread_mutexattr_getrobust_np\(3C\)](#).

A thread that uses mutexes initialized with the `PTHREAD_PRIO_INHERIT` or `PTHREAD_PRIO_PROTECT` *protocol* attribute values should have its *contentionscope* attribute equal to `PTHREAD_SCOPE_SYSTEM` (see [pthread_attr_getscope\(3C\)](#)) and its scheduling policy equal to `SCHED_FIFO` or `SCHED_RR` (see [pthread_attr_getschedparam\(3C\)](#) and [pthread_getschedparam\(3C\)](#)).

If a thread with *contentionscope* attribute equal to `PTHREAD_SCOPE_PROCESS` and/or its scheduling policy equal to `SCHED_OTHER` uses a mutex initialized with the `PTHREAD_PRIO_INHERIT` or `PTHREAD_PRIO_PROTECT` *protocol* attribute value, the effect on the thread's scheduling and priority is unspecified.

The `_POSIX_THREAD_PRIO_INHERIT` and `_POSIX_THREAD_PRIO_PROTECT` options are designed to provide features to solve priority inversion due to mutexes. A priority inheritance or priority ceiling mutex is designed to minimize the dispatch latency of a high priority thread when a low priority thread is holding a mutex required by the high priority thread. This is a specific need for the realtime application domain.

Threads created by realtime applications need to be such that their priorities can influence their access to system resources (CPU resources, at least), in competition with all threads running on the system.

RETURN VALUES

Upon successful completion, the `pthread_mutexattr_getprotocol()` and `pthread_mutexattr_setprotocol()` functions return 0. Otherwise, an error number is returned to indicate the error.

ERRORS

The `pthread_mutexattr_getprotocol()` and `pthread_mutexattr_setprotocol()` functions will fail if:

- | | |
|---------|---|
| EINVAL | The value specified by <i>attr</i> is NULL. |
| ENOSYS | Neither of the options <code>_POSIX_THREAD_PRIO_PROTECT</code> and <code>_POSIX_THREAD_PRIO_INHERIT</code> is defined and the system does not support the function. |
| ENOTSUP | The value specified by <i>protocol</i> is an unsupported value. |

The `pthread_mutexattr_getprotocol()` and `pthread_mutexattr_setprotocol()` functions may fail if:

`pthread_mutexattr_getprotocol(3C)`

`EINVAL` The value specified by *attr* or *protocol* is invalid.
`EPERM` The caller does not have the privilege to perform the operation.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `pthread_create(3C)`, `pthread_mutex_init(3C)`, `pthread_cond_init(3C)`, `pthread_mutexattr_getrobust_np(3C)`, `attributes(5)`, `standards(5)`

pthread_mutexattr_getpshared(3C)

NAME	pthread_mutexattr_getpshared, pthread_mutexattr_setpshared – get and set process-shared attribute
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_mutexattr_getpshared(const pthread_mutexattr_t *restrict attr, int *restrict pshared); int pthread_mutexattr_setpshared(pthread_mutexattr_t *attr, int pshared);</pre>
DESCRIPTION	<p>The <code>pthread_mutexattr_getpshared()</code> function obtains the value of the <i>process-shared</i> attribute from the attributes object referenced by <i>attr</i>. The <code>pthread_mutexattr_setpshared()</code> function is used to set the <i>process-shared</i> attribute in an initialized attributes object referenced by <i>attr</i>.</p> <p>The <i>process-shared</i> attribute is set to <code>PTHREAD_PROCESS_SHARED</code> to permit a mutex to be operated upon by any thread that has access to the memory where the mutex is allocated, even if the mutex is allocated in memory that is shared by multiple processes. If the <i>process-shared</i> attribute is <code>PTHREAD_PROCESS_PRIVATE</code>, the mutex will only be operated upon by threads created within the same process as the thread that initialized the mutex; if threads of differing processes attempt to operate on such a mutex, the behavior is undefined. The default value of the attribute is <code>PTHREAD_PROCESS_PRIVATE</code>.</p>
RETURN VALUES	<p>Upon successful completion, <code>pthread_mutexattr_getpshared()</code> returns 0 and stores the value of the <i>process-shared</i> attribute of <i>attr</i> into the object referenced by the <i>pshared</i> parameter. Otherwise, an error number is returned to indicate the error.</p> <p>Upon successful completion, <code>pthread_mutexattr_setpshared()</code> returns 0. Otherwise, an error number is returned to indicate the error.</p>
ERRORS	<p>The <code>pthread_mutexattr_getpshared()</code> and <code>pthread_mutexattr_setpshared()</code> functions may fail if:</p> <p><code>EINVAL</code> The value specified by <i>attr</i> is invalid.</p> <p>The <code>pthread_mutexattr_setpshared()</code> function may fail if:</p> <p><code>EINVAL</code> The new value specified for the attribute is outside the range of legal values for that attribute.</p>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

`pthread_mutexattr_getpshared(3C)`

SEE ALSO `pthread_create(3C)`, `pthread_mutex_init(3C)`,
`pthread_mutexattr_init(3C)`, `pthread_cond_init(3C)`, `attributes(5)`,
`standards(5)`

pthread_mutexattr_getrobust_np(3C)

NAME	pthread_mutexattr_getrobust_np, pthread_mutexattr_setrobust_np – get or set robustness attribute of mutex attribute object
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutexattr_getrobust_np(const pthread_mutexattr_t *attr, int *robustness); int pthread_mutexattr_setrobust_np(pthread_mutexattr_t *attr, int robustness);</pre>
DESCRIPTION	<p>The following applies only if the symbol <code>_POSIX_THREAD_PRIO_INHERIT</code> is defined, and the mutex attributes object <i>attr</i> should be used only to initialize mutexes that will also be initialized with the protocol attribute having the value <code>PTHREAD_PRIO_INHERIT</code>. See pthread_mutexattr_getprotocol(3C).</p> <p>The <code>pthread_mutexattr_setrobust_np()</code> and <code>pthread_mutexattr_getrobust_np()</code> functions set and get the <i>robustness</i> attribute of a mutex attribute object pointed to by <i>attr</i> that was previously created by the function pthread_mutexattr_init(3C).</p> <p>The <i>robustness</i> attribute defines the behavior when the owner of a mutex dies. The value of <i>robustness</i> may be either <code>PTHREAD_MUTEX_ROBUST_NP</code> or <code>PTHREAD_MUTEX_STALLED_NP</code>, which are defined by the header <code><pthread.h></code>. The default value of the <i>robustness</i> attribute is <code>PTHREAD_MUTEX_STALLED_NP</code>.</p> <p>When the owner of a mutex with the <code>PTHREAD_MUTEX_STALLED_NP</code> <i>robustness</i> attribute dies, all future calls to pthread_mutex_lock(3C) for this mutex will be blocked from progress in an unspecified manner.</p> <p>When the owner of a mutex with the <code>PTHREAD_MUTEX_ROBUST_NP</code> <i>robustness</i> attribute dies, the mutex is unlocked. The next owner of this mutex acquires it with an error value of <code>EOWNERDEAD</code>. Note that the application must always check the return value from <code>pthread_mutex_lock()</code> for a mutex initialized with the <code>PTHREAD_MUTEX_ROBUST_NP</code> <i>robustness</i> attribute. The new owner of this mutex should then attempt to make the state protected by the mutex consistent, since this state could have been left inconsistent when the last owner died. If the new owner is able to make the state consistent, it should call pthread_mutex_consistent_np(3C) for the mutex and then unlock the mutex. If for any reason the new owner is not able to make the state consistent, it should not call <code>pthread_mutex_consistent_np()</code> for the mutex, but should simply unlock the mutex. In the latter scenario, all waiters will be awakened and all subsequent calls to <code>pthread_mutex_lock()</code> will fail in acquiring the mutex with an error value of <code>ENOTRECOVERABLE</code>. The mutex can then be made consistent by uninitialized the mutex with the <code>pthread_mutex_destroy()</code> function and reinitializing it with the <code>pthread_mutex_init()</code> function. If the thread that acquired the lock with <code>EOWNERDEAD</code> dies, the next owner will acquire the lock with an error value of <code>EOWNERDEAD</code>.</p>

pthread_mutexattr_getrobust_np(3C)

Note that the mutex may be in memory shared between processes or in memory private to a process, i.e. the "owner" referenced above is a thread, either within or outside the requestor's process.

The mutex memory must be zeroed before initialization.

RETURN VALUES Upon successful completion, the `pthread_mutexattr_getrobust_np()` and `pthread_mutexattr_setrobust_np()` functions return 0. Otherwise, an error number is returned to indicate the error.

ERRORS The `pthread_mutexattr_getrobust_np()` and `pthread_mutexattr_setrobust_np()` functions will fail if:

- EINVAL** The value specified by *attr* or *robustness* is invalid.
- ENOSYS** The option `_POSIX_THREAD_PRIO_INHERIT` is not defined and the implementation does not support the function.
- ENOTSUP** The value specified by *robustness* is an unsupported value.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `pthread_mutex_lock(3C)`, `pthread_mutex_consistent_np(3C)`, `pthread_mutexattr_getprotocol(3C)`, `attributes(5)`, `mutex(5)`, `standards(5)`

pthread_mutexattr_gettype(3C)

NAME	pthread_mutexattr_gettype, pthread_mutexattr_settype – get or set a mutex type						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutexattr_gettype(pthread_mutexattr_t *restrict <i>attr</i>, int *restrict <i>type</i>); int pthread_mutexattr_settype(pthread_mutexattr_t *<i>attr</i>, int <i>type</i>);</pre>						
DESCRIPTION	<p>The pthread_mutexattr_gettype() and pthread_mutexattr_settype() functions respectively get and set the mutex <i>type</i> attribute. This attribute is set in the <i>type</i> parameter to these functions. The default value of the <i>type</i> attribute is PTHREAD_MUTEX_DEFAULT.</p> <p>The type of mutex is contained in the <i>type</i> attribute of the mutex attributes. Valid mutex types include:</p> <table><tr><td>PTHREAD_MUTEX_NORMAL</td><td>This type of mutex does not detect deadlock. A thread attempting to relock this mutex without first unlocking it will deadlock. Attempting to unlock a mutex locked by a different thread results in undefined behavior. Attempting to unlock an unlocked mutex results in undefined behavior.</td></tr><tr><td>PTHREAD_MUTEX_ERRORCHECK</td><td>This type of mutex provides error checking. A thread attempting to relock this mutex without first unlocking it will return with an error. A thread attempting to unlock a mutex that another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error.</td></tr><tr><td>PTHREAD_MUTEX_RECURSIVE</td><td>A thread attempting to relock this mutex without first unlocking it will succeed in locking the mutex. The relocking deadlock that can occur with mutexes of type PTHREAD_MUTEX_NORMAL cannot occur with this type of mutex. Multiple locks of this mutex require the same number of unlocks to release the mutex before another thread can acquire the mutex. A thread attempting to unlock a mutex that another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error. This type of</td></tr></table>	PTHREAD_MUTEX_NORMAL	This type of mutex does not detect deadlock. A thread attempting to relock this mutex without first unlocking it will deadlock. Attempting to unlock a mutex locked by a different thread results in undefined behavior. Attempting to unlock an unlocked mutex results in undefined behavior.	PTHREAD_MUTEX_ERRORCHECK	This type of mutex provides error checking. A thread attempting to relock this mutex without first unlocking it will return with an error. A thread attempting to unlock a mutex that another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error.	PTHREAD_MUTEX_RECURSIVE	A thread attempting to relock this mutex without first unlocking it will succeed in locking the mutex. The relocking deadlock that can occur with mutexes of type PTHREAD_MUTEX_NORMAL cannot occur with this type of mutex. Multiple locks of this mutex require the same number of unlocks to release the mutex before another thread can acquire the mutex. A thread attempting to unlock a mutex that another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error. This type of
PTHREAD_MUTEX_NORMAL	This type of mutex does not detect deadlock. A thread attempting to relock this mutex without first unlocking it will deadlock. Attempting to unlock a mutex locked by a different thread results in undefined behavior. Attempting to unlock an unlocked mutex results in undefined behavior.						
PTHREAD_MUTEX_ERRORCHECK	This type of mutex provides error checking. A thread attempting to relock this mutex without first unlocking it will return with an error. A thread attempting to unlock a mutex that another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error.						
PTHREAD_MUTEX_RECURSIVE	A thread attempting to relock this mutex without first unlocking it will succeed in locking the mutex. The relocking deadlock that can occur with mutexes of type PTHREAD_MUTEX_NORMAL cannot occur with this type of mutex. Multiple locks of this mutex require the same number of unlocks to release the mutex before another thread can acquire the mutex. A thread attempting to unlock a mutex that another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error. This type of						

pthread_mutexattr_gettype(3C)

mutex is only supported for mutexes whose process shared attribute is PTHREAD_PROCESS_PRIVATE.

PTHREAD_MUTEX_DEFAULT

Attempting to recursively lock a mutex of this type results in undefined behavior. Attempting to unlock a mutex of this type that was not locked by the calling thread results in undefined behavior. Attempting to unlock a mutex of this type that is not locked results in undefined behavior. An implementation is allowed to map this mutex to one of the other mutex types.

RETURN VALUES

Upon successful completion, the pthread_mutexattr_settype() function returns 0. Otherwise, an error number is returned to indicate the error.

Upon successful completion, the pthread_mutexattr_gettype() function returns 0 and stores the value of the type attribute of attr in the object referenced by the type parameter. Otherwise an error number is returned to indicate the error.

ERRORS

The pthread_mutexattr_gettype() and pthread_mutexattr_settype() functions will fail if:

EINVAL The value type is invalid.

The pthread_mutexattr_gettype() and pthread_mutexattr_settype() functions may fail if:

EINVAL The value specified by attr is invalid.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO

pthread_cond_timedwait(3C), pthread_cond_wait(3C), attributes(5), standards(5)

NOTES

Application should not use a PTHREAD_MUTEX_RECURSIVE mutex with condition variables because the implicit unlock performed for pthread_cond_wait() or pthread_cond_timedwait() will not actually release the mutex (if it had been locked multiple times). If this occurs, no other thread can satisfy the condition of the predicate.

pthread_mutexattr_init(3C)

NAME	pthread_mutexattr_init, pthread_mutexattr_destroy – initialize and destroy mutex attributes object						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutexattr_init(pthread_mutexattr_t *attr); int pthread_mutexattr_destroy(pthread_mutexattr_t *attr);</pre>						
DESCRIPTION	<p>The <code>pthread_mutexattr_init()</code> function initializes a mutex attributes object <i>attr</i> with the default value for all of the attributes defined by the implementation.</p> <p>The effect of initializing an already initialized mutex attributes object is undefined.</p> <p>After a mutex attributes object has been used to initialize one or more mutexes, any function affecting the attributes object (including destruction) does not affect any previously initialized mutexes.</p> <p>The <code>pthread_mutexattr_destroy()</code> function destroys a mutex attributes object; the object becomes, in effect, uninitialized. An implementation may cause <code>pthread_mutexattr_destroy()</code> to set the object referenced by <i>attr</i> to an invalid value. A destroyed mutex attributes object can be re-initialized using <code>pthread_mutexattr_init()</code>; the results of otherwise referencing the object after it has been destroyed are undefined.</p>						
RETURN VALUES	Upon successful completion, <code>pthread_mutexattr_init()</code> and <code>pthread_mutexattr_destroy()</code> return 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	The <code>pthread_mutexattr_init()</code> function may fail if: ENOMEM Insufficient memory exists to initialize the mutex attributes object.	The <code>pthread_mutexattr_destroy()</code> function may fail if: EINVAL The value specified by <i>attr</i> is invalid.					
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>pthread_cond_init(3C)</code> , <code>pthread_create(3C)</code> , <code>pthread_mutex_init(3C)</code> , <code>pthread_mutexattr_settype(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	pthread_mutex_consistent_np – make a mutex consistent after owner death
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutex_consistent_np(pthread_mutex_t *<i>mutex</i>);</pre>
DESCRIPTION	<p>The following applies only if the symbol <code>_POSIX_THREAD_PRIO_INHERIT</code> is defined, and for mutexes that have been initialized with the protocol attribute having the value <code>PTHREAD_PRIO_INHERIT</code>. See pthread_mutexattr_getprotocol(3C).</p> <p>The mutex object referenced by <i>mutex</i> is made consistent by calling <code>pthread_mutex_consistent_np()</code>.</p> <p>A consistent mutex becomes inconsistent and is unlocked if its owner dies while holding it. A subsequent owner of the mutex will acquire the mutex with pthread_mutex_lock(3C), which will return <code>EOWNERDEAD</code> to indicate that the acquired mutex is inconsistent.</p> <p>The <code>pthread_mutex_consistent_np()</code> function should be called while holding the mutex acquired by a previous call to <code>pthread_mutex_lock()</code> that returned <code>EOWNERDEAD</code>.</p> <p>Since the critical section protected by the mutex could have been left in an inconsistent state by the dead owner, the caller should make the mutex consistent only if it is able to make the critical section protected by the mutex consistent.</p> <p>Calls to <code>pthread_mutex_lock()</code>, <code>pthread_mutex_unlock()</code>, and <code>pthread_mutex_trylock()</code> for a consistent mutex will behave in the normal manner.</p> <p>The behavior of <code>pthread_mutex_consistent_np()</code> for a mutex which is not inconsistent, or which is not held, is undefined.</p>
RETURN VALUES	Upon successful completion, the <code>pthread_mutexattr_consistent_np()</code> function returns 0. Otherwise, an error number is returned to indicate the error.
ERRORS	<p>The <code>pthread_mutex_consistent_np()</code> function will fail if:</p> <p><code>ENOSYS</code> The option <code>_POSIX_THREAD_PRIO_INHERIT</code> is not defined and the implementation does not support the function.</p> <p>The <code>pthread_mutex_consistent_np()</code> function may fail if:</p> <p><code>EINVAL</code> The value specified by <i>mutex</i> is invalid, or the mutex does not have the appropriate attributes.</p>
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

pthread_mutex_consistent_np(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [pthread_mutex_lock\(3C\)](#), [pthread_mutexattr_getprotocol\(3C\)](#), [pthread_mutexattr_getrobust_np\(3C\)](#), [attributes\(5\)](#), [mutex\(5\)](#), [standards\(5\)](#)

pthread_mutex_getprioceiling(3C)

NAME	pthread_mutex_getprioceiling, pthread_mutex_setprioceiling – change the priority ceiling of a mutex
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutex_getprioceiling(const pthread_mutex_t *restrict mutex, int *restrict prioceiling); int pthread_mutex_setprioceiling(pthread_mutex_t *restrict mutex, int prioceiling, int *restrict old_ceiling);</pre>
DESCRIPTION	<p>The pthread_mutex_getprioceiling() function returns the current priority ceiling of the mutex.</p> <p>The pthread_mutex_setprioceiling() function either locks the mutex if it is unlocked, or blocks until it can successfully lock the mutex, then it changes the mutex's priority ceiling and releases the mutex. When the change is successful, the previous value of the priority ceiling is returned in <i>old_ceiling</i>. The process of locking the mutex need not adhere to the priority protect protocol.</p> <p>If the pthread_mutex_setprioceiling() function fails, the mutex priority ceiling is not changed.</p> <p>The ceiling value should be drawn from the range of priorities for the SCHED_FIFO policy. When a thread acquires such a mutex, the policy of the thread at mutex acquisition should match that from which the ceiling value was derived (SCHED_FIFO, in this case). If a thread changes its scheduling policy while holding a ceiling mutex, the behavior of pthread_mutex_lock() and pthread_mutex_unlock() on this mutex is undefined. See pthread_mutex_lock(3C).</p> <p>The ceiling value should not be treated as a persistent value resident in a pthread_mutex_t that is valid across upgrades of Solaris. The semantics of the actual ceiling value are determined by the existing priority range for the SCHED_FIFO policy, as returned by the sched_get_priority_min() and sched_get_priority_max() functions (see sched_get_priority_min(3RT)) when called on the version of Solaris on which the ceiling value is being utilized.</p>
RETURN VALUES	Upon successful completion, the pthread_mutex_getprioceiling() and pthread_mutex_setprioceiling() functions return 0. Otherwise, an error number is returned to indicate the error.
ERRORS	<p>The pthread_mutex_getprioceiling() and pthread_mutex_setprioceiling() functions will fail if:</p> <p>ENOSYS The option _POSIX_THREAD_PRIO_PROTECT is not defined and the system does not support the function. Since _POSIX_THREAD_PRIO_PROTECT is defined, this condition is not reported.</p> <p>The pthread_mutex_setprioceiling() function will fail if:</p>

pthread_mutex_getprioceiling(3C)

EINVAL The mutex was not initialized with its *protocol* attribute having the value of PTHREAD_PRIO_PROTECT.

The `pthread_mutex_getprioceiling()` and `pthread_mutex_setprioceiling()` functions may fail if:

EINVAL The value specified by *mutex* does not refer to a currently existing mutex.

ENOSYS The system does not support the priority ceiling protocol for mutexes.

EPERM The caller does not have the privilege to perform the operation.

The `pthread_mutex_getprioceiling()` function may fail if:

EINVAL The priority requested by *prioceiling* is out of range.

The `pthread_mutex_setprioceiling()` function may fail if:

EINVAL The priority requested by *prioceiling* is out of range.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO

`pthread_mutex_init(3C)`, `pthread_mutex_lock(3C)`, `sched_get_priority_min(3RT)`, `attributes(5)`, `standards(5)`

NAME	pthread_mutex_init, pthread_mutex_destroy – initialize or destroy a mutex										
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutex_init(pthread_mutex_t *restrict <i>mutex</i>, const pthread_mutexattr_t *restrict <i>attr</i>); int pthread_mutex_destroy(pthread_mutex_t *<i>mutex</i>); pthread_mutex_t <i>mutex</i>= PTHREAD_MUTEX_INITIALIZER</pre>										
DESCRIPTION	<p>The <code>pthread_mutex_init()</code> function initializes the mutex referenced by <i>mutex</i> with attributes specified by <i>attr</i>. If <i>attr</i> is NULL, the default mutex attributes are used; the effect is the same as passing the address of a default mutex attributes object. Upon successful initialization, the state of the mutex becomes initialized and unlocked.</p> <p>Attempting to initialize an already initialized mutex results in undefined behavior.</p> <p>The <code>pthread_mutex_destroy()</code> function destroys the mutex object referenced by <i>mutex</i>; the mutex object becomes, in effect, uninitialized. A destroyed mutex object can be re-initialized using <code>pthread_mutex_init()</code>; the results of otherwise referencing the object after it has been destroyed are undefined.</p> <p>It is safe to destroy an initialized mutex that is unlocked. Attempting to destroy a locked mutex results in undefined behavior.</p> <p>In cases where default mutex attributes are appropriate, the macro <code>PTHREAD_MUTEX_INITIALIZER</code> can be used to initialize mutexes that are statically allocated. The effect is equivalent to dynamic initialization by a call to <code>pthread_mutex_init()</code> with parameter <i>attr</i> specified as NULL, except that no error checks are performed.</p>										
RETURN VALUES	If successful, the <code>pthread_mutex_init()</code> and <code>pthread_mutex_destroy()</code> functions return 0. Otherwise, an error number is returned to indicate the error.										
ERRORS	<p>The <code>pthread_mutex_init()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EAGAIN</td> <td>The system lacked the necessary resources (other than memory) to initialize another mutex.</td> </tr> <tr> <td>EBUSY</td> <td>An attempt was detected to re-initialize the object referenced by <i>mutex</i>, a robust mutex previously initialized but not yet destroyed.</td> </tr> <tr> <td>EINVAL</td> <td>The value specified by <i>attr</i> has not been initialized using <code>pthread_mutexattr_init(3C)</code>.</td> </tr> <tr> <td>EINVAL</td> <td>An attempt was made to initialize <i>mutex</i> to be robust and not <code>PTHREAD_PRIO_INHERIT</code>.</td> </tr> <tr> <td>EINVAL</td> <td>An attempt was made to initialize <i>mutex</i> to be <code>PTHREAD_PRIO_INHERIT</code> or <code>PTHREAD_PRIO_PROTECT</code> and the type of the mutex to be <code>PTHREAD_MUTEX_ERRORCHECK</code> or <code>PTHREAD_MUTEX_RECURSIVE</code>.</td> </tr> </table>	EAGAIN	The system lacked the necessary resources (other than memory) to initialize another mutex.	EBUSY	An attempt was detected to re-initialize the object referenced by <i>mutex</i> , a robust mutex previously initialized but not yet destroyed.	EINVAL	The value specified by <i>attr</i> has not been initialized using <code>pthread_mutexattr_init(3C)</code> .	EINVAL	An attempt was made to initialize <i>mutex</i> to be robust and not <code>PTHREAD_PRIO_INHERIT</code> .	EINVAL	An attempt was made to initialize <i>mutex</i> to be <code>PTHREAD_PRIO_INHERIT</code> or <code>PTHREAD_PRIO_PROTECT</code> and the type of the mutex to be <code>PTHREAD_MUTEX_ERRORCHECK</code> or <code>PTHREAD_MUTEX_RECURSIVE</code> .
EAGAIN	The system lacked the necessary resources (other than memory) to initialize another mutex.										
EBUSY	An attempt was detected to re-initialize the object referenced by <i>mutex</i> , a robust mutex previously initialized but not yet destroyed.										
EINVAL	The value specified by <i>attr</i> has not been initialized using <code>pthread_mutexattr_init(3C)</code> .										
EINVAL	An attempt was made to initialize <i>mutex</i> to be robust and not <code>PTHREAD_PRIO_INHERIT</code> .										
EINVAL	An attempt was made to initialize <i>mutex</i> to be <code>PTHREAD_PRIO_INHERIT</code> or <code>PTHREAD_PRIO_PROTECT</code> and the type of the mutex to be <code>PTHREAD_MUTEX_ERRORCHECK</code> or <code>PTHREAD_MUTEX_RECURSIVE</code> .										

pthread_mutex_init(3C)

- EINVAL An attempt was made to initialize *mutex* to be PTHREAD_PROCESS_SHARED and the type of the mutex to be PTHREAD_MUTEX_RECURSIVE.
- EINVAL The process-shared attribute associated with *attr* is invalid.
- ENOMEM Insufficient memory exists to initialize the mutex.
- EPERM The caller does not have the privilege to perform the operation.

The pthread_mutex_init() function may fail if:

- EBUSY An attempt was detected to re-initialize the object referenced by *mutex*, a mutex previously initialized but not yet destroyed.
- EINVAL The value specified by *attr* or *mutex* is invalid.

The pthread_mutex_destroy() function may fail if:

- EBUSY An attempt was detected to destroy the object referenced by *mutex* while it is locked or referenced (for example, while being used in a pthread_cond_wait(3C) or pthread_cond_timedwait(3C)) by another thread.
- EINVAL The value specified by *mutex* is invalid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO pthread_cond_timedwait(3C), pthread_cond_wait(3C), pthread_mutex_getprioceiling(3C), pthread_mutex_lock(3C), pthread_mutex_unlock(3C), pthread_mutex_setprioceiling(3C), pthread_mutex_trylock(3C), pthread_mutexattr_init(3C), pthread_mutexattr_getpshared(3C), pthread_mutexattr_setpshared(3C) attributes(5), mutex(5), standards(5)

NAME	pthread_mutex_lock, pthread_mutex_trylock, pthread_mutex_unlock – lock or unlock a mutex
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_mutex_lock(pthread_mutex_t *<i>mutex</i>) ; int pthread_mutex_trylock(pthread_mutex_t *<i>mutex</i>) ; int pthread_mutex_unlock(pthread_mutex_t *<i>mutex</i>) ;</pre>
DESCRIPTION	<p>The mutex object referenced by <i>mutex</i> is locked by calling <code>pthread_mutex_lock()</code>. If the mutex is already locked, the calling thread blocks until the mutex becomes available. This operation returns with the mutex object referenced by <i>mutex</i> in the locked state with the calling thread as its owner.</p> <p>If the mutex type is <code>PTHREAD_MUTEX_NORMAL</code>, deadlock detection is not provided. Attempting to relock the mutex causes deadlock. If a thread attempts to unlock a mutex that it has not locked or a mutex that is unlocked, undefined behavior results.</p> <p>If the mutex type is <code>PTHREAD_MUTEX_ERRORCHECK</code>, then error checking is provided. If a thread attempts to relock a mutex that it has already locked, an error will be returned. If a thread attempts to unlock a mutex that it has not locked or a mutex which is unlocked, an error will be returned.</p> <p>If the mutex type is <code>PTHREAD_MUTEX_RECURSIVE</code>, then the mutex maintains the concept of a lock count. When a thread successfully acquires a mutex for the first time, the lock count is set to 1. Every time a thread relocks this mutex, the lock count is incremented by one. Each time the thread unlocks the mutex, the lock count is decremented by one. When the lock count reaches 0, the mutex becomes available for other threads to acquire. If a thread attempts to unlock a mutex that it has not locked or a mutex that is unlocked, an error will be returned.</p> <p>If the mutex type is <code>PTHREAD_MUTEX_DEFAULT</code>, attempting to recursively lock the mutex results in undefined behavior. Attempting to unlock the mutex if it was not locked by the calling thread results in undefined behavior. Attempting to unlock the mutex if it is not locked results in undefined behavior.</p> <p>The <code>pthread_mutex_trylock()</code> function is identical to <code>pthread_mutex_lock()</code> except that if the mutex object referenced by <i>mutex</i> is currently locked (by any thread, including the current thread), the call returns immediately.</p> <p>The <code>pthread_mutex_unlock()</code> function releases the mutex object referenced by <i>mutex</i>. The manner in which a mutex is released is dependent upon the mutex's type attribute. If there are threads blocked on the mutex object referenced by <i>mutex</i> when <code>pthread_mutex_unlock()</code> is called, resulting in the mutex becoming available, the scheduling policy is used to determine which thread will acquire the mutex. (In the case of <code>PTHREAD_MUTEX_RECURSIVE</code> mutexes, the mutex becomes available when the count reaches 0 and the calling thread no longer has any locks on this mutex.)</p>

pthread_mutex_lock(3C)

	<p>If a signal is delivered to a thread waiting for a mutex, upon return from the signal handler the thread resumes waiting for the mutex as if it was not interrupted.</p>
RETURN VALUES	<p>If successful, the <code>pthread_mutex_lock()</code> and <code>pthread_mutex_unlock()</code> functions return 0. Otherwise, an error number is returned to indicate the error.</p> <p>The <code>pthread_mutex_trylock()</code> function returns 0 if a lock on the mutex object referenced by <i>mutex</i> is acquired. Otherwise, an error number is returned to indicate the error.</p>
ERRORS	<p>The <code>pthread_mutex_lock()</code> and <code>pthread_mutex_trylock()</code> functions will fail if:</p> <p>EINVAL The <i>mutex</i> was created with the protocol attribute having the value <code>PTHREAD_PRIO_PROTECT</code> and the calling thread's priority is higher than the mutex's current priority ceiling.</p> <p>The <code>pthread_mutex_trylock()</code> function will fail if:</p> <p>EBUSY The <i>mutex</i> could not be acquired because it was already locked.</p> <p>The <code>pthread_mutex_lock()</code>, <code>pthread_mutex_trylock()</code> and <code>pthread_mutex_unlock()</code> functions may fail if:</p> <p>EINVAL The value specified by <i>mutex</i> does not refer to an initialized mutex object.</p> <p>EAGAIN The mutex could not be acquired because the maximum number of recursive locks for <i>mutex</i> has been exceeded.</p> <p>The <code>pthread_mutex_lock()</code> function may fail if:</p> <p>EDEADLK The current thread already owns the mutex.</p> <p>The <code>pthread_mutex_unlock()</code> function may fail if:</p> <p>EPERM The current thread does not own the mutex.</p> <p>When a thread makes a call to <code>pthread_mutex_lock()</code> or <code>pthread_mutex_trylock()</code>, if the symbol <code>_POSIX_THREAD_PRIO_INHERIT</code> is defined and the mutex is initialized with the protocol attribute having the value <code>PTHREAD_PRIO_INHERIT</code> and the robustness attribute having the value <code>PTHREAD_MUTEX_ROBUST_NP</code> (see <code>pthread_mutexattr_getrobust_np(3C)</code>), the <code>pthread_mutex_lock()</code> and <code>pthread_mutex_trylock()</code> functions will fail if:</p> <p>EOWNERDEAD The last owner of this mutex died while holding the mutex. This mutex is now owned by the caller. The caller must now attempt to make the state protected by the mutex consistent. If it is able to clean up the state, then it should call <code>pthread_mutex_consistent_np()</code> for the mutex and unlock the mutex. Subsequent calls to <code>pthread_mutex_lock()</code> and</p>

pthread_mutex_lock(3C)

pthread_mutex_trylock() will behave normally, as before. If the caller is not able to clean up the state, pthread_mutex_consistent_np() should not be called for the mutex, but it should be unlocked. Subsequent calls to pthread_mutex_lock() and pthread_mutex_trylock() will fail to acquire the mutex with the error value ENOTRECOVERABLE. If the owner who acquired the lock with EOWNERDEAD dies, the next owner will acquire the lock with EOWNERDEAD.

ENOTRECOVERABLE

The mutex trying to be acquired is protecting the state that has been left irrecoverable by the mutex's last owner, who died while holding the lock. The mutex has not been acquired. This condition can occur when the lock was previously acquired with EOWNERDEAD, and the owner was not able to clean up the state and unlocked the mutex without making the mutex consistent.

ENOMEM

The limit on the number of simultaneously held mutexes has been exceeded.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO

pthread_mutex_init(3C), pthread_mutex_destroy(3C), pthread_mutex_consistent_np(3C), pthread_mutexattr_getrobust_np(3C), attributes(5), standards(5)

NOTES

In the current implementation of threads, pthread_mutex_lock(), pthread_mutex_unlock(), mutex_lock(), mutex_unlock(), pthread_mutex_trylock(), and mutex_trylock() do not validate the mutex type. Therefore, an uninitialized mutex or a mutex with an invalid type does not return EINVAL. Interfaces for mutexes with an invalid type have unspecified behavior.

Uninitialized mutexes that are allocated locally may contain junk data. Such mutexes need to be initialized using pthread_mutex_init() or mutex_init().

pthread_mutex_timedlock(3C)

NAME	pthread_mutex_timedlock, pthread_mutex_reltimedlock_np – lock a mutex
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <pthread.h> #include <time.h> int pthread_mutex_timedlock(pthread_mutex_t *restrict mutex, const struct timespec *restrict abs_timeout); int pthread_mutex_reltimedlock_np(pthread_mutex_t *restrict mutex, const struct timespec *restrict rel_timeout);</pre>
DESCRIPTION	<p>The <code>pthread_mutex_timedlock()</code> function locks the mutex object referenced by <i>mutex</i>. If the mutex is already locked, the calling thread blocks until the mutex becomes available as in the <code>pthread_mutex_lock(3C)</code>. If the mutex cannot be locked without waiting for another thread to unlock the mutex, this wait is terminated when the specified timeout expires.</p> <p>The <code>pthread_mutex_reltimedlock_np()</code> function is identical to the <code>pthread_mutex_timedlock()</code> function, except that the timeout is specified as a relative time interval.</p> <p>For <code>pthread_mutex_timedlock()</code>, the timeout expires when the absolute time specified by <i>abs_timeout</i> passes, as measured by the clock on which timeouts are based (that is, when the value of that clock equals or exceeds <i>abs_timeout</i>), or if the absolute time specified by <i>abs_timeout</i> has already been passed at the time of the call.</p> <p>For <code>pthread_mutex_reltimedlock_np()</code>, the timeout expires when the time interval specified by <i>rel_timeout</i> passes, as measured by the <code>CLOCK_REALTIME</code> clock, or if the time interval specified by <i>rel_timeout</i> is negative at the time of the call.</p> <p>The resolution of the timeout is the resolution of the <code>CLOCK_REALTIME</code> clock. The <code>timespec</code> data type is defined in the <code><time.h></code> header.</p> <p>Under no circumstance will either function fail with a timeout if the mutex can be locked immediately. The validity of the timeout parameter need not be checked if the mutex can be locked immediately.</p> <p>As a consequence of the priority inheritance rules (for mutexes initialized with the <code>PRIO_INHERIT</code> protocol), if a timed mutex wait is terminated because its timeout expires, the priority of the owner of the mutex is adjusted as necessary to reflect the fact that this thread is no longer among the threads waiting for the mutex.</p>
RETURN VALUES	Upon successful completion, the <code>pthread_mutex_timedlock()</code> and <code>pthread_mutex_reltimedlock_np()</code> functions return 0. Otherwise, an error number is returned to indicate the error.
ERRORS	The <code>pthread_mutex_timedlock()</code> and <code>pthread_mutex_reltimedlock_np()</code> functions will fail if:
EINVAL	The mutex was created with the protocol attribute having the value <code>PTHREAD_PRIO_PROTECT</code> and the calling thread's priority is higher than the mutex's current priority ceiling.

pthread_mutex_timedlock(3C)

- EINVAL The process or thread would have blocked and the timeout parameter specified a nanoseconds field value less than zero or greater than or equal to 1,000 million.
- ETIMEDOUT The mutex could not be locked before the specified timeout expired.

The pthread_mutex_timedlock() and pthread_mutex_reltimedlock_np() functions may fail if:

- EINVAL The value specified by *mutex* does not refer to an initialized mutex object.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	See below.
MT-Level	MT-Safe

The pthread_mutex_timedlock() is Standard. The pthread_mutex_reltimedlock_np() function is Stable.

SEE ALSO time(2), pthread_mutex_destroy(3C), pthread_mutex_lock(3C), pthread_mutex_trylock(3C), attributes(5), standards(5)

pthread_once(3C)

NAME pthread_once – initialize dynamic package

SYNOPSIS cc -mt [*flag...*] *file...* -lpthread [-lrt *library...*]

```
#include <pthread.h>
pthread_once_t once_control = PTHREAD_ONCE_INIT;

int pthread_once(pthread_once_t *once_control, void
    (*init_routine) (void));
```

DESCRIPTION If any thread in a process with a *once_control* parameter makes a call to `pthread_once()`, the first call will summon the `init_routine()`, but subsequent calls will not. The *once_control* parameter determines whether the associated initialization routine has been called. The `init_routine()` is complete upon return of `pthread_once()`.

`pthread_once()` is not a cancellation point; however, if the function `init_routine()` is a cancellation point and is canceled, the effect on *once_control* is the same as if `pthread_once()` had never been called.

The constant `PTHREAD_ONCE_INIT` is defined in the `<pthread.h>` header.

If *once_control* has automatic storage duration or is not initialized by `PTHREAD_ONCE_INIT`, the behavior of `pthread_once()` is undefined.

RETURN VALUES Upon successful completion, `pthread_once()` returns 0. Otherwise, an error number is returned to indicate the error.

ERRORS `EINVAL` *once_control* or *init_routine* is NULL.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `attributes(5)`, `standards(5)`

NOTES Solaris threads do not offer this functionality.

pthread_rwlockattr_getpshared(3C)

NAME pthread_rwlockattr_getpshared, pthread_rwlockattr_setpshared – get or set process-shared attribute of read-write lock attributes object

SYNOPSIS

```
cc -mt [ flag... ] file... -lpthread [ -lrt library... ]
#include <pthread.h>

int pthread_rwlockattr_getpshared(const pthread_rwlockattr_t
    *restrict attr, int *restrict pshared);

int pthread_rwlockattr_setpshared(pthread_rwlockattr_t *attr, int
    pshared);
```

DESCRIPTION

The *process-shared* attribute is set to PTHREAD_PROCESS_SHARED to permit a read-write lock to be operated upon by any thread that has access to the memory where the read-write lock is allocated, even if the read-write lock is allocated in memory that is shared by multiple processes. If the *process-shared* attribute is PTHREAD_PROCESS_PRIVATE, the read-write lock will only be operated upon by threads created within the same process as the thread that initialised the read-write lock; if threads of differing processes attempt to operate on such a read-write lock, the behaviour is undefined. The default value of the *process-shared* attribute is PTHREAD_PROCESS_PRIVATE.

The pthread_rwlockattr_getpshared() function obtains the value of the *process-shared* attribute from the initialised attributes object referenced by attr. The pthread_rwlockattr_setpshared() function is used to set the *process-shared* attribute in an initialised attributes object referenced by attr.

RETURN VALUES

If successful, the pthread_rwlockattr_setpshared() function returns 0. Otherwise, an error number is returned to indicate the error.

Upon successful completion, the pthread_rwlockattr_getpshared() returns 0 and stores the value of the *process-shared* attribute of attr into the object referenced by the pshared parameter. Otherwise an error number is returned to indicate the error.

ERRORS

The pthread_rwlockattr_getpshared() and pthread_rwlockattr_setpshared() functions will fail if:

EINVAL The value specified by attr or pshared is invalid.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO pthread_rwlock_init(3C), pthread_rwlock_rdlock(3C), pthread_rwlock_unlock(3C), pthread_rwlock_wrlock(3C), pthread_rwlockattr_init(3C), attributes(5), standards(5)

pthread_rwlockattr_init(3C)

NAME	pthread_rwlockattr_init, pthread_rwlockattr_destroy – initialize or destroy read-write lock attributes object						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_rwlockattr_init(pthread_rwlockattr_t *attr); int pthread_rwlockattr_destroy(pthread_rwlockattr_t *attr);</pre>						
DESCRIPTION	<p>The <code>pthread_rwlockattr_init()</code> function initializes a read-write lock attributes object <i>attr</i> with the default value for all of the attributes defined by the implementation.</p> <p>Results are undefined if <code>pthread_rwlockattr_init()</code> is called specifying an already initialized read-write lock attributes object.</p> <p>After a read-write lock attributes object has been used to initialize one or more read-write locks, any function affecting the attributes object (including destruction) does not affect any previously initialized read-write locks.</p> <p>The <code>pthread_rwlockattr_destroy()</code> function destroys a read-write lock attributes object. The effect of subsequent use of the object is undefined until the object is re-initialized by another call to <code>pthread_rwlockattr_init()</code>. An implementation may cause <code>pthread_rwlockattr_destroy()</code> to set the object referenced by <i>attr</i> to an invalid value.</p>						
RETURN VALUES	If successful, the <code>pthread_rwlockattr_init()</code> and <code>pthread_rwlockattr_destroy()</code> functions return 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	<p>The <code>pthread_rwlockattr_init()</code> function will fail if:</p> <table><tr><td>ENOMEM</td><td>Insufficient memory exists to initialize the read-write lock attributes object.</td></tr></table> <p>The <code>pthread_rwlockattr_destroy()</code> function may fail if:</p> <table><tr><td>EINVAL</td><td>The value specified by <i>attr</i> is invalid.</td></tr></table>	ENOMEM	Insufficient memory exists to initialize the read-write lock attributes object.	EINVAL	The value specified by <i>attr</i> is invalid.		
ENOMEM	Insufficient memory exists to initialize the read-write lock attributes object.						
EINVAL	The value specified by <i>attr</i> is invalid.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>pthread_rwlock_init(3C)</code> , <code>pthread_rwlock_rdlock(3C)</code> , <code>pthread_rwlock_unlock(3C)</code> , <code>pthread_rwlock_wrlock(3C)</code> , <code>pthread_rwlockattr_getpshared(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	pthread_rwlock_init, pthread_rwlock_destroy – initialize or destroy a read-write lock object				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -pthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_rwlock_init(pthread_rwlock_t *restrict <i>rwlock</i>, const pthread_rwlockattr_t *restrict <i>attr</i>); int pthread_rwlock_destroy(pthread_rwlock_t **<i>rwlock</i>); pthread_rwlock_t <i>rwlock</i>=PTHREAD_RWLOCK_INITIALIZER;</pre>				
DESCRIPTION	<p>The <code>pthread_rwlock_init()</code> function initializes the read-write lock referenced by <i>rwlock</i> with the attributes referenced by <i>attr</i>. If <i>attr</i> is NULL, the default read-write lock attributes are used; the effect is the same as passing the address of a default read-write lock attributes object. Once initialized, the lock can be used any number of times without being re-initialized. Upon successful initialization, the state of the read-write lock becomes initialized and unlocked. Results are undefined if <code>pthread_rwlock_init()</code> is called specifying an already initialized read-write lock. Results are undefined if a read-write lock is used without first being initialized.</p> <p>If the <code>pthread_rwlock_init()</code> function fails, <i>rwlock</i> is not initialized and the contents of <i>rwlock</i> are undefined.</p> <p>The <code>pthread_rwlock_destroy()</code> function destroys the read-write lock object referenced by <i>rwlock</i> and releases any resources used by the lock. The effect of subsequent use of the lock is undefined until the lock is re-initialized by another call to <code>pthread_rwlock_init()</code>. An implementation may cause <code>pthread_rwlock_destroy()</code> to set the object referenced by <i>rwlock</i> to an invalid value. Results are undefined if <code>pthread_rwlock_destroy()</code> is called when any thread holds <i>rwlock</i>. Attempting to destroy an uninitialized read-write lock results in undefined behaviour. A destroyed read-write lock object can be re-initialized using <code>pthread_rwlock_init()</code>; the results of otherwise referencing the read-write lock object after it has been destroyed are undefined.</p> <p>In cases where default read-write lock attributes are appropriate, the macro <code>PTHREAD_RWLOCK_INITIALIZER</code> can be used to initialize read-write locks that are statically allocated. The effect is equivalent to dynamic initialization by a call to <code>pthread_rwlock_init()</code> with the parameter <i>attr</i> specified as NULL, except that no error checks are performed.</p>				
RETURN VALUES	If successful, the <code>pthread_rwlock_init()</code> and <code>pthread_rwlock_destroy()</code> functions return 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>The <code>pthread_rwlock_init()</code> and <code>pthread_rwlock_destroy()</code> functions will fail if:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">EINVAL</td> <td>The value specified by <i>attr</i> is invalid.</td> </tr> <tr> <td style="padding-right: 20px;">EINVAL</td> <td>The value specified by <i>rwlock</i> is invalid.</td> </tr> </table>	EINVAL	The value specified by <i>attr</i> is invalid.	EINVAL	The value specified by <i>rwlock</i> is invalid.
EINVAL	The value specified by <i>attr</i> is invalid.				
EINVAL	The value specified by <i>rwlock</i> is invalid.				

pthread_rwlock_init(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `pthread_rwlock_rdlock(3C)`, `pthread_rwlock_unlock(3C)`,
`pthread_rwlock_wrlock(3C)`, `pthread_rwlockattr_init(3C)`,
`attributes(5)`, `standards(5)`

NAME	pthread_rwlock_rdlock, pthread_rwlock_tryrdlock – lock or attempt to lock a read-write lock object for reading
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_rwlock_rdlock(pthread_rwlock_t *<i>rwlock</i>); int pthread_rwlock_tryrdlock(pthread_rwlock_t *<i>rwlock</i>);</pre>
DESCRIPTION	<p>The pthread_rwlock_rdlock() function applies a read lock to the read-write lock referenced by <i>rwlock</i>. The calling thread acquires the read lock if a writer does not hold the lock and there are no writers blocked on the lock.</p> <p>The calling thread does not acquire the lock if a writer holds the lock or if writers of higher or equal priority are blocked on the lock; otherwise, the calling thread acquires the lock. If the read lock is not acquired, the calling thread blocks until it can acquire the lock.</p> <p>A thread can hold multiple concurrent read locks on <i>rwlock</i> (that is, successfully call the pthread_rwlock_rdlock() function <i>n</i> times). If so, the thread must perform matching unlocks (that is, it must call the pthread_rwlock_unlock() function <i>n</i> times).</p> <p>The maximum number of concurrent read locks that a thread can hold on one read-write lock is currently set at 100,000, though this number could change in a future release. There is no imposed limit on the number of different threads that can apply a read lock to one read-write lock.</p> <p>The pthread_rwlock_tryrdlock() function applies a read lock like the pthread_rwlock_rdlock() function, with the exception that the function fails if the equivalent pthread_rwlock_rdlock() call would have blocked the calling thread. In no case will the pthread_rwlock_tryrdlock() function ever bloc. It always either acquires the lock or fails and returns immediately.</p> <p>Results are undefined if any of these functions are called with an uninitialized read-write lock.</p> <p>If a signal is delivered to a thread waiting for a read-write lock for reading, upon return from the signal handler the thread resumes waiting for the read-write lock for reading as if it was not interrupted.</p>
RETURN VALUES	<p>If successful, the pthread_rwlock_rdlock() function returns 0. Otherwise, an error number is returned to indicate the error.</p> <p>The pthread_rwlock_tryrdlock() function returns 0 if the lock for reading on the read-write lock object referenced by <i>rwlock</i> is acquired. Otherwise an error number is returned to indicate the error.</p>
ERRORS	The pthread_rwlock_rdlock() and pthread_rwlock_tryrdlock() functions will fail if:

pthread_rwlock_rdlock(3C)

EAGAIN The read lock could not be acquired because the maximum number of read locks by the current thread for *rwlock* has been exceeded.

The `pthread_rwlock_rdlock()` function will fail if:

EDEADLK The current thread already owns the read-write lock for writing.

The `pthread_rwlock_tryrdlock()` function will fail if:

EBUSY The read-write lock could not be acquired for reading because a writer holds the lock or a writer with the appropriate priority was blocked on it.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `pthread_rwlock_init(3C)`, `pthread_rwlock_wrlock(3C)`, `pthread_rwlockattr_init(3C)`, `pthread_rwlock_unlock(3C)`, `attributes(5)`, `standards(5)`

NAME	pthread_rwlock_timedrdlock, pthread_rwlock_reltimedrdlock_np – lock a read-write lock for reading
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <pthread.h> #include <time.h> int pthread_rwlock_timedrdlock(pthread_rwlock_t *restrict rwlock, const struct timespec *restrict abs_timeout); int pthread_rwlock_reltimedrdlock_np(pthread_rwlock_t *restrict rwlock, const struct timespec *restrict rel_timeout);</pre>
DESCRIPTION	<p>The pthread_rwlock_timedrdlock() function applies a read lock to the read-write lock referenced by <i>rwlock</i> as in the pthread_rwlock_rdlock(3C) function. If the lock cannot be acquired without waiting for other threads to unlock the lock, this wait will be terminated when the specified timeout expires. The timeout expires when the absolute time specified by <i>abs_timeout</i> passes, as measured by the CLOCK_REALTIME clock (that is, when the value of that clock equals or exceeds <i>abs_timeout</i>), or if the absolute time specified by <i>abs_timeout</i> has already been passed at the time of the call.</p> <p>The pthread_rwlock_reltimedrdlock_np() function is identical to the pthread_rwlock_timedrdlock() function, except that the timeout is specified as a relative time interval. The timeout expires when the time interval specified by <i>rel_timeout</i> passes, as measured by the CLOCK_REALTIME clock, or if the time interval specified by <i>rel_timeout</i> is negative at the time of the call.</p> <p>The resolution of the timeout is the resolution of the CLOCK_REALTIME clock. The timespec data type is defined in the <time.h> header. Under no circumstances does either function fail with a timeout if the lock can be acquired immediately. The validity of the timeout parameter need not be checked if the lock can be immediately acquired.</p> <p>If a signal that causes a signal handler to be executed is delivered to a thread blocked on a read-write lock with a call to pthread_rwlock_timedrdlock() or pthread_rwlock_reltimedrdlock_np(), upon return from the signal handler the thread resumes waiting for the lock as if it was not interrupted.</p> <p>The calling thread might deadlock if at the time the call is made it holds a write lock on <i>rwlock</i>.</p> <p>The results are undefined if this function is called with an uninitialized read-write lock.</p>
RETURN VALUES	The pthread_rwlock_timedrdlock() and pthread_rwlock_reltimedrdlock_np() functions return 0 if the lock for reading on the read-write lock object referenced by <i>rwlock</i> is acquired. Otherwise, an error number is returned to indicate the error.
ERRORS	The pthread_rwlock_timedrdlock() and pthread_rwlock_reltimedrdlock_np() functions will fail if:

pthread_rwlock_timedrdlock(3C)

ETIMEDOUT The lock could not be acquired before the specified timeout expired.

The `pthread_rwlock_timedrdlock()` and `pthread_rwlock_reltimedrdlock_np()` functions may fail if:

EAGAIN The read lock could not be acquired because the maximum number of read locks for lock would be exceeded.

EDEADLK The calling thread already holds a write lock on *rwlock*.

EINVAL The value specified by *rwlock* does not refer to an initialized read-write lock object, or the timeout nanosecond value is less than zero or greater than or equal to 1 000 million.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	See below.
MT-Level	MT-Safe

The `pthread_rwlock_timedrdlock()` function is Standard. The `pthread_rwlock_reltimedrdlock_np()` is Stable.

SEE ALSO `pthread_rwlock_destroy(3C)`, `pthread_rwlock_rdlock(3C)`,
`pthread_rwlock_timedwrlock(3C)`, `pthread_rwlock_trywrlock(3C)`,
`pthread_rwlock_unlock(3C)`, `pthread_rwlock_wrlock(3C)`, `attributes(5)`,
`standards(5)`

NAME	pthread_rwlock_timedwrlock, pthread_rwlock_reltimedwrlock_np – lock a read-write lock for writing
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <pthread.h> #include <time.h> int pthread_rwlock_timedwrlock(pthread_rwlock_t *restrict rwlock, const struct timespec *restrict abs_timeout); int pthread_rwlock_reltimedwrlock_np(pthread_rwlock_t *restrict rwlock, const struct timespec *restrict rel_timeout);</pre>
DESCRIPTION	<p>The pthread_rwlock_timedwrlock() function applies a write lock to the read-write lock referenced by <i>rwlock</i> as in the pthread_rwlock_wrlock(3C) function. If the lock cannot be acquired without waiting for other threads to unlock the lock, this wait will be terminated when the specified timeout expires. The timeout expires when the absolute time specified by <i>abs_timeout</i> passes, as measured by the CLOCK_REALTIME clock (that is, when the value of that clock equals or exceeds <i>abs_timeout</i>), or if the absolute time specified by <i>abs_timeout</i> has already been passed at the time of the call.</p> <p>The pthread_rwlock_reltimedwrlock_np() function is identical to the pthread_rwlock_timedwrlock() function, except that the timeout is specified as a relative time interval. The timeout expires when the time interval specified by <i>rel_timeout</i> passes, as measured by the CLOCK_REALTIME clock, or if the time interval specified by <i>rel_timeout</i> is negative at the time of the call.</p> <p>The resolution of the timeout is the resolution of the CLOCK_REALTIME clock. The timespec data type is defined in the <time.h> header. Under no circumstances does either function fail with a timeout if the lock can be acquired immediately. The validity of the <i>abs_timeout</i> parameter need not be checked if the lock can be immediately acquired.</p> <p>If a signal that causes a signal handler to be executed is delivered to a thread blocked on a read- write lock with a call to pthread_rwlock_timedwrlock() or pthread_rwlock_reltimedwrlock_np(), upon return from the signal handler the thread resumes waiting for the lock as if it was not interrupted.</p> <p>The calling thread can deadlock if at the time the call is made it holds the read-write lock. The results are undefined if this function is called with an uninitialized read-write lock.</p>
RETURN VALUES	The pthread_rwlock_timedwrlock() and pthread_rwlock_reltimedwrlock_np() functions return 0 if the lock for writing on the read-write lock object referenced by <i>rwlock</i> is acquired. Otherwise, an error number is returned to indicate the error.
ERRORS	The pthread_rwlock_timedwrlock() and pthread_rwlock_reltimedwrlock_np() functions will fail if:

pthread_rwlock_timedwrlock(3C)

ETIMEDOUT The lock could not be acquired before the specified timeout expired.

The `pthread_rwlock_timedwrlock()` and `pthread_rwlock_reltimedwrlock_np()` functions may fail if:

EDEADLK The calling thread already holds the rwlock.

EINVAL The value specified by *rwlock* does not refer to an initialized read-write lock object, or the timeout nanosecond value is less than zero or greater than or equal to 1,000 million.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	See below.
MT-Level	MT-Safe

The `pthread_rwlock_timedwrlock()` function is Standard. The `pthread_rwlock_reltimedwrlock_np()` function is Stable.

SEE ALSO `pthread_rwlock_destroy(3C)`, `pthread_rwlock_rdlock(3C)`, `pthread_rwlock_timedrdlock(3C)`, `pthread_rwlock_trywrlock(3C)`, `pthread_rwlock_unlock(3C)`, `pthread_rwlock_wrlock(3C)`, `attributes(5)`, `standards(5)`

NAME	pthread_rwlock_unlock – unlock a read-write lock object						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -pthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_rwlock_unlock(pthread_rwlock_t *<i>rwlock</i>);</pre>						
DESCRIPTION	<p>The <code>pthread_rwlock_unlock()</code> function is called to release a lock held on the read-write lock object referenced by <code>rwlock</code>. Results are undefined if the read-write lock <code>rwlock</code> is not held by the calling thread.</p> <p>If this function is called to release a read lock from the read-write lock object and there are other read locks currently held on this read-write lock object, the read-write lock object remains in the read locked state. If this function releases the calling thread's last read lock on this read-write lock object, then the calling thread is no longer one of the owners of the object. If this function releases the last read lock for this read-write lock object, the read-write lock object will be put in the unlocked state with no owners.</p> <p>If this function is called to release a write lock for this read-write lock object, the read-write lock object will be put in the unlocked state with no owners.</p> <p>If the call to the <code>pthread_rwlock_unlock()</code> function results in the read-write lock object becoming unlocked and there are multiple threads waiting to acquire the read-write lock object for writing, the scheduling policy is used to determine which thread acquires the read-write lock object for writing. If there are multiple threads waiting to acquire the read-write lock object for reading, the scheduling policy is used to determine the order in which the waiting threads acquire the read-write lock object for reading. If there are multiple threads blocked on <code>rwlock</code> for both read locks and write locks, it is unspecified whether the readers acquire the lock first or whether a writer acquires the lock first.</p> <p>Results are undefined if any of these functions are called with an uninitialized read-write lock.</p>						
RETURN VALUES	If successful, the <code>pthread_rwlock_unlock()</code> function returns 0. Otherwise, an error number is returned to indicate the error.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>pthread_rwlock_init(3C)</code> , <code>pthread_rwlock_rdlock(3C)</code> , <code>pthread_rwlock_wrlock(3C)</code> , <code>pthread_rwlockattr_init(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

pthread_rwlock_wrlock(3C)

NAME	pthread_rwlock_wrlock, pthread_rwlock_trywrlock – lock or attempt to lock a read-write lock object for writing				
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> int pthread_rwlock_wrlock(pthread_rwlock_t *rwlock); int pthread_rwlock_trywrlock(pthread_rwlock_t *rwlock);</pre>				
DESCRIPTION	<p>The pthread_rwlock_wrlock() function applies a write lock to the read-write lock referenced by <i>rwlock</i>. The calling thread acquires the write lock if no other thread (reader or writer) holds the read-write lock <i>rwlock</i>. Otherwise, the thread blocks until it can acquire the lock.</p> <p>The pthread_rwlock_trywrlock() function applies a write lock like the pthread_rwlock_wrlock() function, with the exception that the function fails if any thread currently holds <i>rwlock</i> (for reading or writing).</p> <p>Writers are favored over readers of the same priority to avoid writer starvation. See pthread_rwlock_rdlock(3C).</p> <p>Results are undefined if any of these functions are called with an uninitialized read-write lock.</p> <p>If a signal is delivered to a thread waiting for a read-write lock for writing, upon return from the signal handler the thread resumes waiting for the read-write lock for writing as if it was not interrupted.</p>				
RETURN VALUES	<p>If successful, the pthread_rwlock_wrlock() function returns 0. Otherwise, an error number is returned to indicate the error.</p> <p>The pthread_rwlock_trywrlock() function returns 0 if the lock for writing on the read-write lock object referenced by <i>rwlock</i> is acquired. Otherwise an error number is returned to indicate the error.</p>				
ERRORS	<p>The pthread_rwlock_wrlock() function will fail if:</p> <p>EDEADLK The current thread already owns the read-write lock for writing or reading.</p> <p>The pthread_rwlock_trywrlock() function will fail if:</p> <p>EBUSY The read-write lock could not be acquired for writing because it was already locked for reading or writing.</p>				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center; padding: 5px;">ATTRIBUTE TYPE</th> <th style="text-align: center; padding: 5px;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td style="padding: 5px;">Interface Stability</td> <td style="padding: 5px;">Standard</td> </tr> </tbody> </table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Standard				

pthread_rwlock_wrlock(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [pthread_rwlock_init\(3C\)](#), [pthread_rwlock_unlock\(3C\)](#),
[pthread_rwlockattr_init\(3C\)](#), [pthread_rwlock_rdlock\(3C\)](#),
[attributes\(5\)](#), [standards\(5\)](#)

pthread_self(3C)

NAME pthread_self – get calling thread’s ID

SYNOPSIS `cc -mt [flag...] file... -lpthread [-lrt library...]`

```
#include <pthread.h>
pthread_t pthread_self(void);
```

DESCRIPTION The `pthread_self()` function returns the thread ID of the calling thread.

ERRORS No errors are defined.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [pthread_create\(3C\)](#), [pthread_equal\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	pthread_setcancelstate – enable or disable cancellation
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_setcancelstate(int <i>state</i>, int *<i>oldstate</i>);</pre>
DESCRIPTION	<p>pthread_setcancelstate() atomically sets the calling thread's cancellation state to the specified <i>state</i> and if <i>oldstate</i> is not NULL, stores the previous cancellation <i>state</i> in <i>oldstate</i>.</p> <p>The <i>state</i> can be either of the following:</p> <p>PTHREAD_CANCEL_ENABLE This is the default. When cancellation is deferred (deferred cancellation is also the default), cancellation occurs when the target thread reaches a cancellation point and a cancel is pending. When cancellation is asynchronous, receipt of a pthread_cancel(3C) call causes immediate cancellation.</p> <p>PTHREAD_CANCEL_DISABLE When cancellation is deferred, all cancellation requests to the target thread are held pending. When cancellation is asynchronous, all cancellation requests to the target thread are held pending; as soon as cancellation is re-enabled, pending cancellations are executed immediately.</p> <p>See cancellation(5) for the definition of a cancellation point and a discussion of cancellation concepts. See pthread_setcanceltype(3C) for explanations of deferred and asynchronous cancellation.</p> <p>pthread_setcancelstate() is a cancellation point when it is called with PTHREAD_CANCEL_ENABLE and the cancellation type is PTHREAD_CANCEL_ASYNCHRONOUS.</p>
RETURN VALUES	Upon successful completion, pthread_setcancelstate(), returns 0. Otherwise, an error number is returned to indicate the error.
ERRORS	<p>The pthread_setcancelstate() function will fail if:</p> <p>EINVAL The specified <i>state</i> is not PTHREAD_CANCEL_ENABLE or PTHREAD_CANCEL_DISABLE.</p>
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

pthread_setcancelstate(3C)

SEE ALSO | pthread_cancel(3C), pthread_cleanup_pop(3C),
pthread_cleanup_push(3C), pthread_exit(3C), pthread_join(3C),
pthread_setcanceltype(3C), pthread_testcancel(3C), setjmp(3C),
attributes(5), cancellation(5), condition(5), standards(5)

NAME	pthread_setcanceltype – set the cancellation type of a thread				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_setcanceltype(int <i>type</i>, int *<i>oldtype</i>);</pre>				
DESCRIPTION	<p>pthread_setcanceltype() atomically sets the calling thread's cancellation type to the specified <i>type</i> and, if <i>oldtype</i> is not NULL, stores the previous cancellation type in <i>oldtype</i>. The <i>type</i> can be either of the following:</p> <table border="0"> <tr> <td style="vertical-align: top;">PTHREAD_CANCEL_DEFERRED</td> <td>This is the default. When cancellation is enabled (enabled cancellation is also the default), cancellation occurs when the target thread reaches a cancellation point and a cancel is pending. When cancellation is disabled, all cancellation requests to the target thread are held pending.</td> </tr> <tr> <td style="vertical-align: top;">PTHREAD_CANCEL_ASYNCHEOUS</td> <td>When cancellation is enabled, receipt of a pthread_cancel(3C) call causes immediate cancellation. When cancellation is disabled, all cancellation requests to the target thread are held pending; as soon as cancellation is re-enabled, pending cancellations are executed immediately.</td> </tr> </table> <p>See cancellation(5) for the definition of a cancellation point and a discussion of cancellation concepts. See pthread_setcancelstate(3C) for explanations of enabling and disabling cancellation.</p> <p>pthread_setcanceltype() is a cancellation point if <i>type</i> is called with PTHREAD_CANCEL_ASYNCHEOUS and the cancellation state is PTHREAD_CANCEL_ENABLE.</p>	PTHREAD_CANCEL_DEFERRED	This is the default. When cancellation is enabled (enabled cancellation is also the default), cancellation occurs when the target thread reaches a cancellation point and a cancel is pending. When cancellation is disabled, all cancellation requests to the target thread are held pending.	PTHREAD_CANCEL_ASYNCHEOUS	When cancellation is enabled, receipt of a pthread_cancel(3C) call causes immediate cancellation. When cancellation is disabled, all cancellation requests to the target thread are held pending; as soon as cancellation is re-enabled, pending cancellations are executed immediately.
PTHREAD_CANCEL_DEFERRED	This is the default. When cancellation is enabled (enabled cancellation is also the default), cancellation occurs when the target thread reaches a cancellation point and a cancel is pending. When cancellation is disabled, all cancellation requests to the target thread are held pending.				
PTHREAD_CANCEL_ASYNCHEOUS	When cancellation is enabled, receipt of a pthread_cancel(3C) call causes immediate cancellation. When cancellation is disabled, all cancellation requests to the target thread are held pending; as soon as cancellation is re-enabled, pending cancellations are executed immediately.				
RETURN VALUES	Upon successful completion, the pthread_setcanceltype() function returns 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>The pthread_setcanceltype() function will fail if:</p> <table border="0"> <tr> <td style="vertical-align: top;">EINVAL</td> <td>The specified <i>type</i> is not PTHREAD_CANCEL_DEFERRED or PTHREAD_CANCEL_ASYNCHEOUS.</td> </tr> </table>	EINVAL	The specified <i>type</i> is not PTHREAD_CANCEL_DEFERRED or PTHREAD_CANCEL_ASYNCHEOUS.		
EINVAL	The specified <i>type</i> is not PTHREAD_CANCEL_DEFERRED or PTHREAD_CANCEL_ASYNCHEOUS.				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

pthread_setcanceltype(3C)

SEE ALSO | pthread_cancel(3C), pthread_cleanup_pop(3C),
pthread_cleanup_push(3C), pthread_exit(3C), pthread_join(3C),
pthread_setcancelstate(3C), pthread_testcancel(3C), setjmp(3C),
attributes(5), cancellation(5), condition(5), standards(5)

NAME	pthread_setschedprio – dynamic thread scheduling parameters access										
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> int pthread_setschedprio(pthread_t <i>thread</i>, int <i>prio</i>);</pre>										
DESCRIPTION	<p>The <code>pthread_setschedprio()</code> function sets the scheduling priority for the thread whose thread ID is given by <i>thread</i> to the value given by <i>prio</i>.</p> <p>If the <code>pthread_setschedprio()</code> function fails, the scheduling priority of the target thread is not changed.</p>										
RETURN VALUES	If successful, the <code>pthread_setschedprio()</code> function returns 0; otherwise, an error number is returned to indicate the error.										
ERRORS	<p>The <code>pthread_setschedprio()</code> function will fail if:</p> <table border="0"> <tr> <td style="vertical-align: top;">EINVAL</td> <td>The value of <i>prio</i> is invalid for the scheduling policy of the specified thread.</td> </tr> <tr> <td style="vertical-align: top;">ENOTSUP</td> <td>An attempt was made to set the priority to an unsupported value.</td> </tr> <tr> <td style="vertical-align: top;">EPERM</td> <td>The caller does not have the appropriate permission to set the scheduling policy of the specified thread.</td> </tr> <tr> <td></td> <td>The system does not allow the application to modify the priority to the value specified.</td> </tr> <tr> <td style="vertical-align: top;">ESRCH</td> <td>The value specified by <i>thread</i> does not refer to an existing thread.</td> </tr> </table>	EINVAL	The value of <i>prio</i> is invalid for the scheduling policy of the specified thread.	ENOTSUP	An attempt was made to set the priority to an unsupported value.	EPERM	The caller does not have the appropriate permission to set the scheduling policy of the specified thread.		The system does not allow the application to modify the priority to the value specified.	ESRCH	The value specified by <i>thread</i> does not refer to an existing thread.
EINVAL	The value of <i>prio</i> is invalid for the scheduling policy of the specified thread.										
ENOTSUP	An attempt was made to set the priority to an unsupported value.										
EPERM	The caller does not have the appropriate permission to set the scheduling policy of the specified thread.										
	The system does not allow the application to modify the priority to the value specified.										
ESRCH	The value specified by <i>thread</i> does not refer to an existing thread.										
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:										
	<table border="1"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe				
ATTRIBUTE TYPE	ATTRIBUTE VALUE										
Interface Stability	Standard										
MT-Level	MT-Safe										
SEE ALSO	<code>pthread_getschedparam(3C)</code> , <code>attributes(5)</code>										

pthread_sigmask(3C)

NAME	pthread_sigmask – change or examine calling thread’s signal mask						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> -lpthread [-lrt <i>library...</i>] #include <pthread.h> #include <signal.h> int pthread_sigmask(int <i>how</i>, const sigset_t *<i>set</i>, sigset_t *<i>oset</i>);</pre>						
DESCRIPTION	<p>The <code>pthread_sigmask()</code> function changes or examines a calling thread’s signal mask. Each thread has its own signal mask. A new thread inherits the calling thread’s signal mask and priority; however, pending signals are not inherited. Signals pending for a new thread will be empty.</p> <p>If the value of the argument <i>set</i> is not NULL, <i>set</i> points to a set of signals that can modify the currently blocked set. If the value of <i>set</i> is NULL, the value of <i>how</i> is insignificant and the thread’s signal mask is unmodified; thus, <code>pthread_sigmask()</code> can be used to inquire about the currently blocked signals.</p> <p>The value of the argument <i>how</i> specifies the method in which the set is changed and takes one of the following values:</p> <table><tr><td>SIG_BLOCK</td><td><i>set</i> corresponds to a set of signals to block. They are added to the current signal mask.</td></tr><tr><td>SIG_UNBLOCK</td><td><i>set</i> corresponds to a set of signals to unblock. These signals are deleted from the current signal mask.</td></tr><tr><td>SIG_SETMASK</td><td><i>set</i> corresponds to the new signal mask. The current signal mask is replaced by <i>set</i>.</td></tr></table> <p>If the value of <i>oset</i> is not NULL, it points to the location where the previous signal mask is stored.</p>	SIG_BLOCK	<i>set</i> corresponds to a set of signals to block. They are added to the current signal mask.	SIG_UNBLOCK	<i>set</i> corresponds to a set of signals to unblock. These signals are deleted from the current signal mask.	SIG_SETMASK	<i>set</i> corresponds to the new signal mask. The current signal mask is replaced by <i>set</i> .
SIG_BLOCK	<i>set</i> corresponds to a set of signals to block. They are added to the current signal mask.						
SIG_UNBLOCK	<i>set</i> corresponds to a set of signals to unblock. These signals are deleted from the current signal mask.						
SIG_SETMASK	<i>set</i> corresponds to the new signal mask. The current signal mask is replaced by <i>set</i> .						
RETURN VALUES	Upon successful completion, the <code>pthread_sigmask()</code> function returns 0. Otherwise, it returns a non-zero value.						
ERRORS	The <code>pthread_sigmask()</code> function will fail if: EINVAL The value of <i>how</i> is not defined and <i>oset</i> is NULL.						
EXAMPLES	<p>EXAMPLE 1 Create a default thread that can serve as a signal catcher/handler with its own signal mask.</p> <p>The following example shows how to create a default thread that can serve as a signal catcher/handler with its own signal mask. <i>new</i> will have a different value from the creator’s signal mask.</p> <p>As POSIX threads and Solaris threads are fully compatible even within the same process, this example uses <code>pthread_create(3C)</code> if you execute <code>a.out 0</code>, or <code>thr_create(3C)</code> if you execute <code>a.out 1</code>.</p> <p>In this example:</p>						

EXAMPLE 1 Create a default thread that can serve as a signal catcher/handler with its own signal mask. (Continued)

- The `sigemptyset(3C)` function initializes a null signal set, `new`. The `sigaddset(3C)` function packs the signal, `SIGINT`, into that new set.
- Either `pthread_sigmask()` or `thr_sigsetmask()` is used to mask the signal, `SIGINT` (CTRL-C), from the calling thread, which is `main()`. The signal is masked to guarantee that only the new thread will receive this signal.
- `pthread_create()` or `thr_create()` creates the signal-handling thread.
- Using `pthread_join(3C)` or `thr_join(3C)`, `main()` then waits for the termination of that signal-handling thread, whose ID number is `user_threadID`; `main()` will then `sleep(3C)` for 2 seconds, after which the program terminates.
- The signal-handling thread, handler:
 - Assigns the handler `interrupt()` to handle the signal `SIGINT`, by the call to `sigaction(2)`.
 - Resets its own signal set to *not block* the signal, `SIGINT`.
 - Sleeps for 8 seconds to allow time for the user to deliver the signal, `SIGINT`, by pressing the CTRL-C.

```

/* cc thisfile.c -lthread -lpthread */
#define _REENTRANT /* basic first 3-lines for threads */
#include <pthread.h>
#include <thread.h>
thread_t user_threadID;
sigset_t new;
void *handler( ), interrupt( );

main( int argc, char *argv[ ] ) {
    test_argv(argv[1]);

    sigemptyset(&new);
    sigaddset(&new, SIGINT);
    switch(*argv[1]) {

        case '0': /* POSIX */
            pthread_sigmask(SIG_BLOCK, &new, NULL);
            pthread_create(&user_threadID, NULL, handler, argv[1]);
            pthread_join(user_threadID, NULL);
            break;

        case '1': /* Solaris */
            thr_sigsetmask(SIG_BLOCK, &new, NULL);
            thr_create(NULL, 0, handler, argv[1], 0, &user_threadID);
            thr_join(user_threadID, NULL, NULL);
            break;
    } /* switch */

    printf("thread handler, # %d, has exited\n",user_threadID);
    sleep(2);
    printf("main thread, # %d is done\n", thr_self( ));
}

```

pthread_sigmask(3C)

EXAMPLE 1 Create a default thread that can serve as a signal catcher/handler with its own signal mask. (Continued)

```
    } /* end main */

    struct sigaction act;

    void *
    handler(char argv1[ ])
    {
        act.sa_handler = interrupt;
        sigaction(SIGINT, &act, NULL);
        switch(*argv1)    {
            case '0':    /* POSIX */
                pthread_sigmask(SIG_UNBLOCK, &new, NULL);
                break;
            case '1':    /* Solaris */
                thr_sigsetmask(SIG_UNBLOCK, &new, NULL);
                break;
        }
        printf("\n Press CTRL-C to deliver SIGINT signal to the process\n");
        sleep(8); /* give user time to hit CTRL-C */
    }

    void
    interrupt(int sig)
    {
        printf("thread %d caught signal %d\n", thr_self( ), sig);
    }

    void test_argv(char argv1[ ])    {
        if(argv1 == NULL)    {
            printf("use 0 as arg1 to use thr_create( );\n \
or use 1 as arg1 to use pthread_create( )\n");
            exit(NULL);
        }
    }
}
```

In the last example, the handler thread served as a signal-handler while also taking care of activity of its own (in this case, sleeping, although it could have been some other activity). A thread could be completely dedicated to signal-handling simply by waiting for the delivery of a selected signal by blocking with `sigwait(2)`. The two subroutines in the previous example, `handler()` and `interrupt()`, could have been replaced with the following routine:

```
void *
handler( )
{
    int signal;
    printf("thread %d is waiting for you to press the CTRL-C keys\n",
        thr_self( ));
    sigwait(&new, &signal);
    printf("thread %d has received the signal %d \n", thr_self( ), signal);
}
/* pthread_create( ) and thr_create( ) would use NULL instead of argv[1]
```

EXAMPLE 1 Create a default thread that can serve as a signal catcher/handler with its own signal mask. (Continued)

```
for the arg passed to handler( ) */
```

In this routine, one thread is dedicated to catching and handling the signal specified by the set new, which allows main() and all of its other sub-threads, created *after* pthread_sigmask() or thr_sigsetmask() masked that signal, to continue uninterrupted. Any use of sigwait(2) should be such that all threads block the signals passed to sigwait(2) at all times. Only the thread that calls sigwait() will get the signals. The call to sigwait(2) takes two arguments.

For this type of background dedicated signal-handling routine, a Solaris daemon thread can be used by passing the argument THR_DAEMON to thr_create(3C).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe and Async-Signal-Safe

SEE ALSO sigaction(2), sigprocmask(2), sigwait(2), cond_wait(3C), pthread_create(3C), pthread_join(3C), pthread_self(3C), sigaddset(3C), sigemptyset(3C), sigsetops(3C), sleep(3C), attributes(5), standards(5)

NOTES It is not possible to block signals that cannot be ignored (see sigaction(2)). If using the threads library, it is not possible to block or unblock the signals SIGWAITING, SIGLWP, or SIGCANCEL. This restriction is quietly enforced by the threads library.

Using sigwait(2) in a dedicated thread allows asynchronously generated signals to be managed synchronously; however, sigwait(2) should never be used to manage synchronously generated signals.

Synchronously generated signals are exceptions that are generated by a thread and are directed at the thread causing the exception. Since sigwait() blocks waiting for signals, the blocking thread cannot receive a synchronously generated signal.

If sigprocmask(2) is used in a multithreaded application, it behaves the same as if pthread_sigmask() has been called. POSIX leaves the semantics of the call to sigprocmask(2) unspecified in a multi-threaded process, so programs that care about POSIX portability should not depend on this semantic.

If a signal is delivered while a thread is waiting on a condition variable, the cond_wait(3C) function will be interrupted and the handler will be executed. The state of the lock protecting the condition variable is undefined while the thread is executing the signal handler.

Although pthread_sigmask() is Async-Signal-Safe with respect to the Solaris environment, this safeness is not guaranteed to be portable to other POSIX domains.

pthread_sigmask(3C)

Signals that are generated synchronously should not be masked. If such a signal is blocked and delivered, the receiving process is killed.

NAME	pthread_spin_destroy, pthread_spin_init – destroy or initialize a spin lock object
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <pthread.h> int pthread_spin_destroy(pthread_spinlock_t *lock); int pthread_spin_init(pthread_spinlock_t *lock, int pshared);</pre>
DESCRIPTION	<p>The <code>pthread_spin_destroy()</code> function destroys the spin lock referenced by <i>lock</i> and release any resources used by the lock. The effect of subsequent use of the lock is undefined until the lock is reinitialized by another call to <code>pthread_spin_init()</code>. The results are undefined if <code>pthread_spin_destroy()</code> is called when a thread holds the lock, or if this function is called with an uninitialized thread spin lock.</p> <p>The <code>pthread_spin_init()</code> function allocates any resources required to use the spin lock referenced by <i>lock</i> and initialize the lock to an unlocked state.</p> <p>If the Thread Process-Shared Synchronization option is supported and the value of <i>pshared</i> is <code>PTHREAD_PROCESS_SHARED</code>, the spin lock can be operated upon by any thread that has access to the memory where the spin lock is allocated, even if it is allocated in memory that is shared by multiple processes.</p> <p>If the Thread Process-Shared Synchronization option is supported and the value of <i>pshared</i> is <code>PTHREAD_PROCESS_PRIVATE</code>, or if the option is not supported, the spin lock can only be operated upon by threads created within the same process as the thread that initialized the spin lock. If threads of differing processes attempt to operate on such a spin lock, the behavior is undefined.</p> <p>The results are undefined if <code>pthread_spin_init()</code> is called specifying an already initialized spin lock. The results are undefined if a spin lock is used without first being initialized.</p> <p>If the <code>pthread_spin_init()</code> function fails, the lock is not initialized and the contents of <i>lock</i> are undefined.</p> <p>Only the object referenced by <i>lock</i> can be used for performing synchronization.</p> <p>The result of referring to copies of that object in calls to <code>pthread_spin_destroy()</code>, <code>pthread_spin_lock(3C)</code>, <code>pthread_spin_trylock(3C)</code>, or <code>pthread_spin_unlock(3C)</code> is undefined.</p>
RETURN VALUES	Upon successful completion, these functions returns 0. Otherwise, an error number is returned to indicate the error.
ERRORS	<p>The <code>pthread_spin_init()</code> function will fail if:</p> <p>EAGAIN The system lacks the necessary resources to initialize another spin lock.</p> <p>These functions may fail if:</p>

pthread_spin_destroy(3C)

EBUSY The system has detected an attempt to initialize or destroy a spin lock while it is in use (for example, while being used in a `pthread_spin_lock()` call) by another thread.

EINVAL The value specified by *lock* is invalid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [pthread_spin_lock\(3C\)](#), [pthread_spin_unlock\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	pthread_spin_lock, pthread_spin_trylock – lock a spin lock object						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> [<i>library...</i>] #include <pthread.h> int pthread_spin_lock(pthread_spinlock_t *<i>lock</i>) ; #include <pthread.h> int pthread_spin_trylock(pthread_spinlock_t *<i>lock</i>) ;</pre>						
DESCRIPTION	<p>The <code>pthread_spin_lock()</code> function locks the spin lock referenced by <i>lock</i>. The calling thread acquires the lock if it is not held by another thread. Otherwise, the thread spins (that is, does not return from the <code>pthread_spin_lock</code> call()) until the lock becomes available. The results are undefined if the calling thread holds the lock at the time the call is made.</p> <p>The <code>pthread_spin_trylock()</code> function locks the spin lock referenced by <i>lock</i> if it is not held by any thread. Otherwise, the function fails.</p> <p>The results are undefined if either of these functions is called with an uninitialized spin lock.</p>						
RETURN VALUES	Upon successful completion, these functions returns 0. Otherwise, an error number is returned to indicate the error.						
ERRORS	<p>The <code>pthread_spin_trylock()</code> function will fail if:</p> <p>EBUSY A thread currently holds the lock.</p> <p>These functions may fail if:</p> <p>EINVAL The value specified by <i>lock</i> does not refer to an initialized spin lock object.</p> <p>The <code>pthread_spin_lock()</code> function may fail if:</p> <p>EDEADLK The calling thread already holds the lock.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>pthread_spin_destroy(3C)</code> , <code>pthread_spin_unlock(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

pthread_spin_unlock(3C)

NAME pthread_spin_unlock – unlock a spin lock object

SYNOPSIS

```
cc -mt [ flag... ] file... [ library... ]
#include <pthread.h>

int pthread_spin_unlock(pthread_spinlock_t *lock);
```

DESCRIPTION The pthread_spin_unlock() function releases the spin lock referenced by *lock* which was locked with the pthread_spin_lock(3C) or pthread_spin_trylock(3C) functions. The results are undefined if the lock is not held by the calling thread. If there are threads spinning on the lock when pthread_spin_unlock() is called, the lock becomes available and an unspecified spinning thread acquires the lock.

The results are undefined if this function is called with an uninitialized thread spin lock.

RETURN VALUES Upon successful completion, the pthread_spin_unlock() function returns 0. Otherwise, an error number shall be returned to indicate the error.

ERRORS The pthread_spin_unlock() function will fail if:

EINVAL	An invalid argument was specified.
EPERM	The calling thread does not hold the lock.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO pthread_spin_destroy(3C), pthread_spin_lock(3C), attributes(5), standards(5)

NAME	pthread_testcancel – create cancellation point in the calling thread
SYNOPSIS	<pre>cc -mt [flag...] file... -lpthread [-lrt library...] #include <pthread.h> void pthread_testcancel(void);</pre>
DESCRIPTION	The pthread_testcancel() function forces testing for cancellation. This is useful when you need to execute code that runs for long periods without encountering cancellation points; such as a library routine that executes long-running computations without cancellation points. This type of code can block cancellation for unacceptable long periods of time. One strategy for avoiding blocking cancellation for long periods, is to insert calls to pthread_testcancel() in the long-running computation code and to setup a cancellation handler in the library code, if required.
RETURN VALUES	The pthread_testcancel() function returns a void.
ERRORS	The pthread_testcancel() function does not return errors.
EXAMPLES	See cancellation(5) for an example of using pthread_testcancel() to force testing for cancellation and a discussion of cancellation concepts.
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO Intro(3), pthread_cleanup_pop(3C), pthread_cleanup_push(3C), pthread_exit(3C), pthread_join(3C), pthread_setcancelstate(3C), pthread_setcanceltype(3C), setjmp(3C), attributes(5), cancellation(5), condition(5), standards(5)

NOTES pthread_testcancel() has no effect if cancellation is disabled.

Use pthread_testcancel() with pthread_setcanceltype() called with its canceltype set to PTHREAD_CANCEL_DEFERRED. The pthread_testcancel() function operation is undefined if pthread_setcanceltype() was called with its canceltype argument set to PTHREAD_CANCEL_ASYNCHRONOUS.

It is possible to kill a thread when it is holding a resource, such as lock or allocated memory. If that thread has not setup a cancellation cleanup handler to release the held resource, the application is "cancel-unsafe". See attributes(5) for a discussion of Cancel-Safety, Deferred-Cancel-Safety, and Asynchronous-Cancel-Safety.

ptrace(3C)

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 <code>ptrace()</code> 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 <code>signal.h(3HEAD)</code>), at which time it enters a stopped state and its parent is notified by the <code>wait(3C)</code> function. When the child is in the stopped state, its parent can examine and modify its “core image” using <code>ptrace()</code>. 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 <code>ptrace()</code> 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 <code>signal(3C)</code>). 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 <code>-1</code> 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 <code>-1</code> 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 these functions, it stops before executing the first instruction of the new image showing signal `SIGTRAP`.

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

- `EIO` The *request* argument is an illegal number.
- `EPERM` The calling process does not have appropriate privileges to control the calling process. See `proc(4)`.
- `ESRCH` The *pid* argument identifies a child that does not exist or has not executed a `ptrace()` call with request 0.

USAGE The `ptrace()` function is available only with the 32-bit version of `libc(3LIB)`. It is not available with the 64-bit version of this library.

The `/proc` debugging interfaces should be used instead of `ptrace()`, which provides quite limited debugger support and is itself implemented using the `/proc` interfaces. There is no actual `ptrace()` system call in the kernel. See `proc(4)` for descriptions of the `/proc` debugging interfaces.

ptrace(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `exec(2)`, `exit(2)`, `libc(3LIB)`, `signal(3C)`, `signal.h(3HEAD)`, `wait(3C)`, `proc(4)`, `attributes(5)`

NAME	ptsname – get name of the slave pseudo-terminal device						
SYNOPSIS	<pre>#include <stdlib.h> char *ptsname(int fildes);</pre>						
DESCRIPTION	The ptsname() function returns the name of the slave pseudo-terminal device associated with a master pseudo-terminal device. <i>fildes</i> is a file descriptor returned from a successful open of the master device. ptsname() returns a pointer to a string containing the null-terminated path name of the slave device of the form /dev/pts/N, where N is a non-negative integer.						
RETURN VALUES	Upon successful completion, the function ptsname() returns a pointer to a string which is the name of the pseudo-terminal slave device. This value points to a static data area that is overwritten by each call to ptsname(). Upon failure, ptsname() returns NULL. This could occur if <i>fildes</i> is an invalid file descriptor or if the slave device name does not exist in the file system.						
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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	open(2), grantpt(3C), ttyname(3C), unlockpt(3C), attributes(5), standards(5) <i>STREAMS Programming Guide</i>						

putenv(3C)

NAME	putenv – change or add value to environment						
SYNOPSIS	<pre>#include <stdlib.h> int putenv(char *string);</pre>						
DESCRIPTION	<p>The <code>putenv()</code> function makes the value of the environment variable <i>name</i> equal to <i>value</i> by altering an existing variable or creating a new one. In either case, the string pointed to by <i>string</i> becomes part of the environment, so altering the string will change the environment.</p> <p>The <i>string</i> argument points to a string of the form <i>name=value</i>. The space used by <i>string</i> is no longer used once a new string-defining <i>name</i> is passed to <code>putenv()</code>.</p> <p>The <code>putenv()</code> function uses <code>malloc(3C)</code> to enlarge the environment.</p> <p><code>putenv()</code> is called, environment variables are not in alphabetical order.</p>						
RETURN VALUES	Upon successful completion, <code>putenv()</code> returns 0. Otherwise, it returns a non-zero value and sets <code>errno</code> to indicate the error.						
ERRORS	The <code>putenv()</code> function may fail if: ENOMEM Insufficient memory was available.						
USAGE	The <code>putenv()</code> function can be safely called from multithreaded programs. Caution must be exercised when using this function and <code>getenv(3C)</code> in multithreaded programs. These functions examine and modify the environment list, which is shared by all threads in a program. The system prevents the list from being accessed simultaneously by two different threads. It does not, however, prevent two threads from successively accessing the environment list using <code>putenv()</code> or <code>getenv()</code> .						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	<code>exec(2)</code> , <code>getenv(3C)</code> , <code>malloc(3C)</code> , <code>attributes(5)</code> , <code>environ(5)</code> , <code>standards(5)</code>						
WARNINGS	The <i>string</i> argument should not be an automatic variable. It should be declared static if it is declared within a function because it cannot be automatically declared. A potential error is to call <code>putenv()</code> with a pointer to an automatic variable as the argument and to then exit the calling function while <i>string</i> is still part of the environment.						

- NAME** | putpwent – write password file entry
- SYNOPSIS** | `#include <pwd.h>`
`int putpwent(const struct passwd *p, FILE *f);`
- DESCRIPTION** | The `putpwent()` function is the inverse of `getpwent()`. See [getpwnam\(3C\)](#). Given a pointer to a `passwd` structure created by `getpwent()`, `getpwuid()`, or `getpwnam()`, `putpwent()` writes a line on the stream `f` that matches the format of `/etc/passwd`.
- RETURN VALUES** | The `putpwent()` function returns a non-zero value if an error was detected during its operation. Otherwise, it returns 0.
- USAGE** | The `putpwent()` function is of limited utility, since most password files are maintained as Network Information Service (NIS) files that cannot be updated with this function. For this reason, the use of this function is discouraged. If used at all, it should be used with [putspent\(3C\)](#) to update the shadow file.
- ATTRIBUTES** | See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Unsafe

SEE ALSO | [getpwnam\(3C\)](#), [putspent\(3C\)](#), [attributes\(5\)](#)

puts(3C)

NAME	puts, fputs – put a string on a stream						
SYNOPSIS	<pre>#include <stdio.h> int puts(const char *s); int fputs(const char *s, FILE *stream);</pre>						
DESCRIPTION	<p>The <code>puts()</code> function writes the string pointed to by <code>s</code>, followed by a <code>NEWLINE</code> character, to the standard output stream <code>stdout</code> (see intro(3)). The terminating null byte is not written.</p> <p>The <code>fputs()</code> function writes the null-terminated string pointed to by <code>s</code> to the named output <code>stream</code>. The terminating null byte is not written.</p> <p>The <code>st_ctime</code> and <code>st_mtime</code> fields of the file will be marked for update between the successful execution of <code>fputs()</code> and the next successful completion of a call to fflush(3C) or fclose(3C) on the same stream or a call to exit(2) or abort(3C).</p>						
RETURN VALUES	On successful completion, both functions return the number of bytes written; otherwise they return <code>EOF</code> and set <code>errno</code> to indicate the error.						
ERRORS	Refer to fputc(3C) .						
USAGE	Unlike <code>puts()</code> , the <code>fputs()</code> function does not write a <code>NEWLINE</code> character at the end of the string.						
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>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	exit(2) , write(2) , intro(3) , abort(3C) , fclose(3C) , ferror(3C) , fflush(3C) , fopen(3C) , fputc(3C) , printf(3C) , stdio(3C) , attributes(5) , standards(5)						

NAME	putspent – write shadow password file entry				
SYNOPSIS	<pre>#include <shadow.h> int putspent(const struct spwd *p, FILE *fp);</pre>				
DESCRIPTION	<p>The <code>putspent()</code> function is the inverse of <code>getspent()</code>. See getspnam(3C). Given a pointer to a <code>spwd</code> structure created by <code>getspent()</code> or <code>getspnam()</code>, <code>putspent()</code> writes a line on the stream <code>fp</code> that matches the format of <code>/etc/shadow</code>.</p> <p>The <code>spwd</code> structure contains the following members:</p> <pre>char *sp_namp; char *sp_pwdp; long sp_lstchg; long sp_min; long sp_max; long sp_warn; long sp_inact; long sp_expire; unsigned long sp_flag;</pre> <p>If the <code>sp_min</code>, <code>sp_max</code>, <code>sp_lstchg</code>, <code>sp_warn</code>, <code>sp_inact</code>, or <code>sp_expire</code> member of the <code>spwd</code> structure is <code>-1</code>, or if <code>sp_flag</code> is <code>0</code>, the corresponding <code>/etc/shadow</code> field is cleared.</p>				
RETURN VALUES	The <code>putspent()</code> function returns a non-zero value if an error was detected during its operation. Otherwise, it returns <code>0</code> .				
USAGE	Since this function is for internal use only, compatibility is not guaranteed. For this reason, its use is discouraged. If used at all, it should be used with putpwent(3C) to update the password file.				
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>Unsafe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Unsafe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Unsafe				
SEE ALSO	getpwnam(3C) , getspnam(3C) , putpwent(3C) , attributes(5)				

putws(3C)

NAME putws – convert a string of Process Code characters to EUC characters

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

`int putws(wchar_t *s);`

DESCRIPTION The `putws()` function converts the Process Code string (terminated by a `(wchar_t) NULL`) pointed to by `s`, to an Extended Unix Code (EUC) string followed by a NEWLINE character, and writes it to the standard output stream `stdout`. It does not write the terminal null character.

RETURN VALUES The `putws()` function returns the number of Process Code characters transformed and written. It returns EOF if it attempts to write to a file that has not been opened for writing.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `ferror(3C)`, `fopen(3C)`, `fread(3C)`, `getws(3C)`, `printf(3C)`, `putwc(3C)`, `attributes(5)`

NAME	qsort – quick sort
SYNOPSIS	<pre>#include <stdlib.h> void qsort(void *base, size_t nel, size_t width, int (*compar)(const void *, const void *));</pre>
DESCRIPTION	<p>The <code>qsort()</code> function is an implementation of the quick-sort algorithm. It sorts a table of data in place. The contents of the table are sorted in ascending order according to the user-supplied comparison function.</p> <p>The <i>base</i> argument points to the element at the base of the table. The <i>nel</i> argument is the number of elements in the table. The <i>width</i> argument specifies the size of each element in bytes. The <i>compar</i> argument is the name of the comparison function, which is called with two arguments that point to the elements being compared.</p> <p>The function must return an integer less than, equal to, or greater than zero to indicate if the first argument is to be considered less than, equal to, or greater than the second argument.</p> <p>The contents of the table are sorted in ascending order according to the user supplied comparison function.</p>
USAGE	The <code>qsort()</code> function safely allows concurrent access by multiple threads to disjoint data, such as overlapping subtrees or tables.
EXAMPLES	<p>EXAMPLE 1 Program sorts.</p> <p>The following program sorts a simple array:</p> <pre>#include <stdlib.h> #include <stdio.h> static int intcompare(const void *p1, const void *p2) { int i = *((int *)p1); int j = *((int *)p2); if (i > j) return (1); if (i < j) return (-1); return (0); } int main() { int i; int a[10] = { 9, 8, 7, 6, 5, 4, 3, 2, 1, 0 }; size_t nelems = sizeof (a) / sizeof (int); qsort((void *)a, nelems, sizeof (int), intcompare);</pre>

qsort(3C)

EXAMPLE 1 Program sorts. *(Continued)*

```
for (i = 0; i < nelems; i++) {  
    (void) printf("%d ", a[i]);  
}  
  
(void) printf("\n");  
return (0);  
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `sort(1)`, `bsearch(3C)`, `lsearch(3C)`, `string(3C)`, `attributes(5)`, `standards(5)`

NOTES The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

The relative order in the output of two items that compare as equal is unpredictable.

NAME	raise – send signal to program						
SYNOPSIS	<pre>#include <signal.h> int raise(int sig);</pre>						
DESCRIPTION	<p>The <code>raise()</code> function sends the signal <code>sig</code> to the executing program. It uses the <code>kill()</code> function to send the signal to the executing program, as follows:</p> <pre>kill(getpid(), sig);</pre> <p>See the <code>kill(2)</code> manual page for a detailed list of failure conditions and the <code>signal(3C)</code> manual page for a list of signals.</p>						
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>getpid(2)</code> , <code>kill(2)</code> , <code>signal(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

rand(3C)

NAME	rand, srand, rand_r – simple random-number generator						
SYNOPSIS	<pre>#include <stdlib.h> int rand(void); void srand(unsigned int seed); int rand_r(unsigned int *seed);</pre>						
DESCRIPTION	<p>The <code>rand()</code> function uses a multiplicative congruential random-number generator with period 2^{32} that returns successive pseudo-random numbers in the range of 0 to <code>RAND_MAX</code> (defined in <code><stdlib.h></code>).</p> <p>The <code>srand()</code> function uses the argument <code>seed</code> as a seed for a new sequence of pseudo-random numbers to be returned by subsequent calls to <code>rand()</code>. If <code>srand()</code> is then called with the same <code>seed</code> value, the sequence of pseudo-random numbers will be repeated. If <code>rand()</code> is called before any calls to <code>srand()</code> have been made, the same sequence will be generated as when <code>srand()</code> is first called with a <code>seed</code> value of 1.</p> <p>The <code>rand_r()</code> function has the same functionality as <code>rand()</code> except that a pointer to a seed <code>seed</code> must be supplied by the caller. If <code>rand_r()</code> is called with the same initial value for the object pointed to by <code>seed</code> and that object is not modified between successive calls to <code>rand_r()</code>, the same sequence as that produced by calls to <code>rand()</code> will be generated.</p> <p>The <code>rand()</code> and <code>srand()</code> functions provide per-process pseudo-random streams shared by all threads. The same effect can be achieved if all threads call <code>rand_r()</code> with a pointer to the same seed object. The <code>rand_r()</code> function allows a thread to generate a private pseudo-random stream by having the seed object be private to the thread.</p>						
USAGE	<p>The spectral properties of <code>rand()</code> are limited. The drand48(3C) function provides a better, more elaborate random-number generator.</p> <p>When compiling multithreaded applications, the <code>_REENTRANT</code> flag must be defined on the compile line. This flag should be used only in multithreaded applications.</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>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	drand48(3C) , attributes(5) , standards(5)						

NAME	rand, srand – simple random number generator
SYNOPSIS	<pre><i>/usr/ucb/cc</i> [<i>flag ...</i>] <i>file ...</i> int rand() int srand(<i>seed</i>); unsigned <i>seed</i>;</pre>
DESCRIPTION	<p>rand() uses a multiplicative congruential random number generator with period 2^{32} to return successive pseudo-random numbers in the range from 0 to $2^{31} - 1$.</p> <p>srand() can be called at any time to reset the random-number generator to a random starting point. The generator is initially seeded with a value of 1.</p>
SEE ALSO	drand48(3C) , rand(3C) , random(3C)
NOTES	<p>Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.</p> <p>The spectral properties of rand() leave a great deal to be desired. drand48(3C) and random(3C) provide much better, though more elaborate, random-number generators.</p> <p>The low bits of the numbers generated are not very random; use the middle bits. In particular the lowest bit alternates between 0 and 1.</p>

random(3C)

NAME	random, srandom, initstate, setstate – pseudorandom number functions
SYNOPSIS	<pre>#include <stdlib.h> long random(void); void srandom(unsigned int seed); char *initstate(unsigned int seed, char *state, size_t size); char *setstate(const char *state);</pre>
DESCRIPTION	<p>The <code>random()</code> function uses a nonlinear additive feedback random-number generator employing a default state array size of 31 long integers to return successive pseudo-random numbers in the range from 0 to $2^{31}-1$. The period of this random-number generator is approximately $16 \times (2^{31}-1)$. The size of the state array determines the period of the random-number generator. Increasing the state array size increases the period.</p> <p>The <code>srandom()</code> function initializes the current state array using the value of <i>seed</i>.</p> <p>The <code>random()</code> and <code>srandom()</code> functions have (almost) the same calling sequence and initialization properties as <code>rand()</code> and <code>srand()</code> (see rand(3C)). The difference is that rand(3C) produces a much less random sequence—in fact, the low dozen bits generated by <code>rand</code> go through a cyclic pattern. All the bits generated by <code>random()</code> are usable.</p> <p>The algorithm from <code>rand()</code> is used by <code>srandom()</code> to generate the 31 state integers. Because of this, different <code>srandom()</code> seeds often produce, within an offset, the same sequence of low order bits from <code>random()</code>. If low order bits are used directly, <code>random()</code> should be initialized with <code>setstate()</code> using high quality random values.</p> <p>Unlike <code>srand()</code>, <code>srandom()</code> does not return the old seed because the amount of state information used is much more than a single word. Two other routines are provided to deal with restarting/changing random number generators. With 256 bytes of state information, the period of the random-number generator is greater than 2^{69}, which should be sufficient for most purposes.</p> <p>Like rand(3C), <code>random()</code> produces by default a sequence of numbers that can be duplicated by calling <code>srandom()</code> with 1 as the seed.</p> <p>The <code>initstate()</code> and <code>setstate()</code> functions handle restarting and changing random-number generators. The <code>initstate()</code> function allows a state array, pointed to by the <i>state</i> argument, to be initialized for future use. The <i>size</i> argument, which specifies the size in bytes of the state array, is used by <code>initstate()</code> to decide what type of random-number generator to use; the larger the state array, the more random the numbers. Values for the amount of state information are 8, 32, 64, 128, and 256 bytes. Other values greater than 8 bytes are rounded down to the nearest one of these values. For values smaller than 8, <code>random()</code> uses a simple linear congruential random number generator. The <i>seed</i> argument specifies a starting point for the random-number sequence and provides for restarting at the same point. The <code>initstate()</code> function returns a pointer to the previous state information array.</p>

If `initstate()` has not been called, then `random()` behaves as though `initstate()` had been called with `seed = 1` and `size = 128`.

If `initstate()` is called with `size < 8`, then `random()` uses a simple linear congruential random number generator.

Once a state has been initialized, `setstate()` allows switching between state arrays. The array defined by the `state` argument is used for further random-number generation until `initstate()` is called or `setstate()` is called again. The `setstate()` function returns a pointer to the previous state array.

RETURN VALUES

The `random()` function returns the generated pseudo-random number.

The `srandom()` function returns no value.

Upon successful completion, `initstate()` and `setstate()` return a pointer to the previous state array. Otherwise, a null pointer is returned.

ERRORS

No errors are defined.

USAGE

After initialization, a state array can be restarted at a different point in one of two ways:

- The `initstate()` function can be used, with the desired seed, state array, and size of the array.
- The `setstate()` function, with the desired state, can be used, followed by `srandom()` with the desired seed. The advantage of using both of these functions is that the size of the state array does not have to be saved once it is initialized.

EXAMPLES

EXAMPLE 1 Initialize an array.

The following example demonstrates the use of `initstate()` to initialize an array. It also demonstrates how to initialize an array and pass it to `setstate()`.

```
# include <stdlib.h>
static unsigned int state0[32];
static unsigned int state1[32] = {
    3,
    0x9a319039, 0x32d9c024, 0x9b663182, 0x5da1f342,
    0x7449e56b, 0xbef1dbb0, 0xab5c5918, 0x946554fd,
    0x8c2e680f, 0xeb3d799f, 0xb11ee0b7, 0x2d436b86,
    0xda672e2a, 0x1588ca88, 0xe369735d, 0x904f35f7,
    0xd7158fd6, 0x6fa6f051, 0x616e6b96, 0xac94efdc,
    0xde3b81e0, 0xdf0a6fb5, 0xf103bc02, 0x48f340fb,
    0x36413f93, 0xc622c298, 0xf5a42ab8, 0x8a88d77b,
    0xf5ad9d0e, 0x8999220b, 0x27fb47b9
};
main() {
    unsigned seed;
    int n;
    seed = 1;
    n = 128;
    (void)initstate(seed, (char *)state0, n);
    printf("random() = %d0\n", random());
}
```

random(3C)

EXAMPLE 1 Initialize an array. (Continued)

```
(void) setstate((char *)state1);
printf("random() = %d0\n", random());
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	See NOTES below.

SEE ALSO [drand48\(3C\)](#), [rand\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NOTES The `random()` and `srandom()` functions are unsafe in multithreaded applications.

Use of these functions in multithreaded applications is unsupported.

For `initstate()` and `setstate()`, the *state* argument must be aligned on an `int` boundary.

Newer and better performing random number generators such as `addrans()` and `lcrans()` are available with the SUNWsprow package.

NAME	rctlblk_set_value, rctlblk_get_firing_time, rctlblk_get_global_action, rctlblk_get_global_flags, rctlblk_get_local_action, rctlblk_get_local_flags, rctlblk_get_privilege, rctlblk_get_recipient_pid, rctlblk_get_value, rctlblk_get_enforced_value, rctlblk_set_local_action, rctlblk_set_local_flags, rctlblk_set_privilege, rctlblk_set_recipient_pid, rctlblk_size – manipulate resource control blocks
SYNOPSIS	<pre>#include <rctl.h> hrtime_t rctlblk_get_firing_time(rctlblk_t *rblk); int rctlblk_get_global_action(rctlblk_t *rblk); int rctlblk_get_global_flags(rctlblk_t *rblk); int rctlblk_get_local_action(rctlblk_t *rblk, int *signalp); int rctlblk_get_local_flags(rctlblk_t *rblk); rctl_priv_t rctlblk_get_privilege(rctlblk_t *rblk); id_t rctlblk_get_recipient_pid(rctlblk_t *rblk); rctl_qty_t rctlblk_get_value(rctlblk_t *rblk); rctl_qty_t rctlblk_get_enforced_value(rctlblk_t *rblk); void rctlblk_set_local_action(rctlblk_t *rblk, rctl_action_t action, int signal); void rctlblk_set_local_flags(rctlblk_t *rblk, int flags); void rctlblk_set_privilege(rctlblk_t *rblk, rctl_priv_t privilege); void rctlblk_set_value(rctlblk_t *rblk, rctl_qty_t value); void rctlblk_set_recipient_pid(id_t pid); size_t rctlblk_size(void);</pre>
DESCRIPTION	<p>The resource control block routines allow the establishment or retrieval of values from a resource control block used to transfer information using the <code>getrctl(2)</code> and <code>setrctl(2)</code> functions. Each of the routines accesses or sets the resource control block member corresponding to its name. Certain of these members are read-only and do not possess set routines.</p> <p>The firing time of a resource control block is 0 if the resource control action-value has not been exceeded for its lifetime on the process. Otherwise the firing time is the value of <code>gethrtime(3C)</code> at the moment the action on the resource control value was taken.</p> <p>The global actions and flags are the action and flags set by <code>rctladm(1M)</code>. These values cannot be set with <code>setrctl(2)</code>. Valid global actions are listed in the table below. Global flags are generally a published property of the control and are not modifiable.</p>

rctlblk_set_value(3C)

RCTL_GLOBAL_DENY_ALWAYS

The action taken when a control value is exceeded on this control will always include denial of the resource.

RCTL_GLOBAL_DENY_NEVER

The action taken when a control value is exceeded on this control will always exclude denial of the resource; the resource will always be granted, although other actions can also be taken.

RCTL_GLOBAL_SIGNAL_NEVER

No signal actions are permitted on this control.

RCTL_GLOBAL_CPU_TIME

The valid signals available as local actions include the SIGXCPU signal.

RCTL_GLOBAL_FILE_SIZE

The valid signals available as local actions include the SIGXFSZ signal.

RCTL_GLOBAL_INFINITE

This resource control supports the concept of an unlimited value; generally true only of accumulation-oriented resources, such as CPU time.

RCTL_GLOBAL_LOWERABLE

Non-privileged callers are able to lower the value of privileged resource control values on this control.

RCTL_GLOBAL_NOACTION

No global action will be taken when a resource control value is exceeded on this control.

RCTL_GLOBAL_NOBASIC

No values with the RCPRIV_BASIC privilege are permitted on this control.

RCTL_GLOBAL_SYSLOG

A standard message will be logged by the [syslog\(3C\)](#) facility when any resource control value on a sequence associated with this control is exceeded.

RCTL_GLOBAL_UNOBSERVABLE

The resource control (generally on a task- or project-related control) does not support observational control values. An RCPRIV_BASIC privileged control value placed by a process on the task or process will generate an action only if the value is exceeded by that process.

RCTL_GLOBAL_BYTES

This resource control represents a number of bytes.

RCTL_GLOBAL_SECONDS

This resource control represents a quantity of time in seconds.

RCTL_GLOBAL_COUNT

This resource control represents an integer count.

The local action and flags are those on the current resource control value represented by this resource control block. Valid actions and flags are listed in the table below. In the case of RCTL_LOCAL_SIGNAL, the second argument to

`rctlblk_set_local_action()` contains the signal to be sent. Similarly, the signal to be sent is copied into the integer location specified by the second argument to `rctlblk_get_local_action()`. A restricted set of signals is made available for normal use by the resource control facility: `SIGBART`, `SIGXRES`, `SIGHUP`, `SIGSTOP`, `SIGTERM`, and `SIGKILL`. Other signals are permitted due to global properties of a specific control. Calls to `setrctl()` with illegal signals will fail.

`RCTL_LOCAL_DENY`

When this resource control value is encountered, the request for the resource will be denied. Set on all values if `RCTL_GLOBAL_DENY_ALWAYS` is set for this control; cleared on all values if `RCTL_GLOBAL_DENY_NEVER` is set for this control.

`RCTL_LOCAL_MAXIMAL`

This resource control value represents a request for the maximum amount of resource for this control. If `RCTL_GLOBAL_INFINITE` is set for this resource control, `RCTL_LOCAL_MAXIMAL` indicates an unlimited resource control value, one that will never be exceeded.

`RCTL_LOCAL_NOACTION`

No local action will be taken when this resource control value is exceeded.

`RCTL_LOCAL_SIGNAL`

The specified signal, sent by `rctlblk_set_local_action()`, will be sent to the process that placed this resource control value in the value sequence. This behavior is also true for signal actions on project and task resource controls. The specified signal is sent only to the recipient process, not all processes within the project or task.

The `rctlblk_get_recipient_pid()` function returns the value of the process ID that placed the resource control value for basic rctls. For privileged or system rctls, `rctlblk_get_recipient_pid()` returns -1.

The `rctlblk_set_recipient_pid()` function sets the recipient *pid* for a basic rctl. When `setrctl(2)` is called with the flag `RCTL_USE_RECIPIENT_PID`, this *pid* is used. Otherwise, the PID of the calling process is used. Only privileged users can set the recipient PID to one other than the PID of the calling process. Process-scoped rctls must have a recipient PID that matches the PID of the calling process.

The `rctlblk_get_privilege()` function returns the privilege of the resource control block. Valid privileges are `RCPRIV_BASIC`, `RCPRIV_PRIVILEGED`, and `RCPRIV_SYSTEM`. System resource controls are read-only. Privileged resource controls require the `{PRIV_SYS_RESOURCE}` privilege to write, unless the `RCTL_GLOBAL_LOWERABLE` global flag is set, in which case unprivileged applications can lower the value of a privileged control.

rctlblk_set_value(3C)

The `rctlblk_get_value()` and `rctlblk_set_value()` functions return or establish the enforced value associated with the resource control. In cases where the process, task, or project associated with the control possesses fewer capabilities than allowable by the current value, the value returned by `rctlblk_get_enforced_value()` will differ from that returned by `rctlblk_get_value()`. This capability difference arises with processes using an address space model smaller than the maximum address space model supported by the system.

The `rctlblk_size()` function returns the size of a resource control block for use in memory allocation. The `rctlblk_t *` type is an opaque pointer whose size is not connected with that of the resource control block itself. Use of `rctlblk_size()` is illustrated in the example below.

RETURN VALUES The various set routines have no return values. Incorrectly composed resource control blocks will generate errors when used with `setrctl(2)` or `getrctl(2)`.

ERRORS No error values are returned. Incorrectly constructed resource control blocks will be rejected by the system calls.

EXAMPLES **EXAMPLE 1** Display the contents of a fetched resource control block.

The following example displays the contents of a fetched resource control block.

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

rctlblk_t *rblk;
int rsignal;
int raction;

if ((rblk = malloc(rctlblk_size())) == NULL) {
    (void) perror("rblk malloc");
    exit(1);
}

if (getrctl("process.max-cpu-time", NULL, rblk, RCTL_FIRST) == -1) {
    (void) perror("getrctl");
    exit(1);
}

main()
{
    raction = rctlblk_get_local_action(rblk, &rsignal),
    (void) printf("Resource control for %s\n",
        "process.max-cpu-time");
    (void) printf("Process ID:      %d\n",
        rctlblk_get_recipient_pid(rblk));
    (void) printf("Privilege:      %x\n"
        rctlblk_get_privilege(rblk));
    (void) printf("Global flags:   %x\n"
        rctlblk_get_global_flags(rblk));
    (void) printf("Global actions: %x\n"
        rctlblk_get_global_action(rblk));
```


EXAMPLE 1 Display the contents of a fetched resource control block. *(Continued)*

```
(void) printf("Local flags:   %x\n"
             rctlblk_get_local_flags(rblk));
(void) printf("Local action:  %x (%d)\n"
             raction, raction == RCTL_LOCAL_SIGNAL ? rsignal : 0);
(void) printf("Value:        %llu\n",
             rctlblk_get_value(rblk));
(void) printf("\t\tEnforced value: %llu\n",
             rctlblk_get_enforced_value(rblk));
}
```

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving
MT-Level	MT-Safe

SEE ALSO [rctladm\(1M\)](#), [getrctl\(2\)](#), [setrctl\(2\)](#), [gethrtime\(3C\)](#), [attributes\(5\)](#)

rctl_walk(3C)

NAME	rctl_walk – visit registered rctls on current system
SYNOPSIS	<pre>#include <rctl.h> int rctl_walk(int (*callback) (const char *rctlname, void *walk_data), void *init_data);</pre>
DESCRIPTION	The rctl_walk() function provides a mechanism for the application author to examine all active resource controls (rctls) on the current system. The <i>callback</i> function provided by the application is given the name of an rctl at each invocation and can use the <i>walk_data</i> to record its own state. The callback function should return non-zero if it encounters an error condition or attempts to terminate the walk prematurely; otherwise the callback function should return 0.
RETURN VALUES	Upon successful completion, rctl_walk() returns 0. It returns -1 if the <i>callback</i> function returned a non-zero value or if the walk encountered an error, in which case <i>errno</i> is set to indicate the error.
ERRORS	The rctl_walk() function will fail if: ENOMEM There is insufficient memory available to set up the initial data for the walk. Other returned error values are presumably caused by the <i>callback</i> function.
EXAMPLES	<p>EXAMPLE 1 Count the number of rctls available on the system.</p> <p>The following example counts the number of resource controls on the system.</p> <pre>#include <sys/types.h> #include <rctl.h> #include <stdio.h> typedef struct wdata { uint_t count; } wdata_t; wdata_t total_count; int simple_callback(const char *name, void *pvt) { wdata_t *w = (wdata_t *)pvt; w->count++; return (0); } ... total_count.count = 0; errno = 0; if (rctl_walk(simple_callback, &total_count) == 0) (void) printf("count = %u\n", total_count.count);</pre>

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving
MT-Level	MT-Safe

SEE ALSO `setrctl(2)`, `attributes(5)`

readdir(3C)

NAME	readdir, readdir_r – read directory
SYNOPSIS	<pre>#include <sys/types.h> #include <dirent.h> struct dirent *readdir(DIR *dirp); struct dirent *readdir_r(DIR *dirp, struct dirent *entry);</pre>
Standard conforming	<pre>cc [flag ...] file ... -D_POSIX_PTHREAD_SEMANTICS [library ...] int readdir_r(DIR *restrict dirp, struct dirent *restrict entry, struct dirent **restrict result);</pre>
DESCRIPTION	<p>The type <code>DIR</code>, which is defined in the header <code><dirent.h></code>, represents a <i>directory stream</i>, which is an ordered sequence of all the directory entries in a particular directory. Directory entries represent files; files may be removed from a directory or added to a directory asynchronously to the operation of <code>readdir()</code> and <code>readdir_r()</code>.</p>
readdir()	<p>The <code>readdir()</code> function returns a pointer to a structure representing the directory entry at the current position in the directory stream specified by the argument <i>dirp</i>, and positions the directory stream at the next entry. It returns a null pointer upon reaching the end of the directory stream. The structure <code>dirent</code> defined by the <code><dirent.h></code> header describes a directory entry.</p> <p>The <code>readdir()</code> function will not return directory entries containing empty names. If entries for <code>.</code> (dot) or <code>..</code> (dot-dot) exist, one entry will be returned for dot and one entry will be returned for dot-dot; otherwise they will not be returned.</p> <p>The pointer returned by <code>readdir()</code> points to data which may be overwritten by another call to <code>readdir()</code> on the same directory stream. This data is not overwritten by another call to <code>readdir()</code> on a different directory stream.</p> <p>If a file is removed from or added to the directory after the most recent call to <code>opendir(3C)</code> or <code>rewinddir(3C)</code>, whether a subsequent call to <code>readdir()</code> returns an entry for that file is unspecified.</p> <p>The <code>readdir()</code> function may buffer several directory entries per actual read operation; <code>readdir()</code> marks for update the <i>st_atime</i> field of the directory each time the directory is actually read.</p> <p>After a call to <code>fork(2)</code>, either the parent or child (but not both) may continue processing the directory stream using <code>readdir()</code>, <code>rewinddir()</code> or <code>seekdir(3C)</code>. If both the parent and child processes use these functions, the result is undefined.</p> <p>If the entry names a symbolic link, the value of the <code>d_ino</code> member is unspecified.</p>
readdir_r()	<p>Unless the end of the directory stream has been reached or an error occurred, the <code>readdir_r()</code> function initializes the <code>dirent</code> structure referenced by <i>entry</i> to represent the directory entry at the current position in the directory stream referred to by <i>dirp</i>, and positions the directory stream at the next entry.</p>

The caller must allocate storage pointed to by *entry* to be large enough for a `dirent` structure with an array of `char d_name` member containing at least `NAME_MAX` (that is, `pathconf(directory, _PC_NAME_MAX)`) plus one elements. (`_PC_NAME_MAX` is defined in `<unistd.h>`.)

The `readdir_r()` function will not return directory entries containing empty names. It is unspecified whether entries are returned for `.` (dot) or `..` (dot-dot).

If a file is removed from or added to the directory after the most recent call to `opendir()` or `rewinddir()`, whether a subsequent call to `readdir_r()` returns an entry for that file is unspecified.

The `readdir_r()` function may buffer several directory entries per actual read operation; the `readdir_r()` function marks for update the `st_atime` field of the directory each time the directory is actually read.

The standard-conforming version (see `standards(5)`) of the `readdir_r()` function performs all of the actions described above and sets the pointer pointed to by *result*. If a directory entry is returned, the pointer will be set to the same value as the *entry* argument; otherwise, it will be set to `NULL`.

RETURN VALUES

Upon successful completion, `readdir()` and the default `readdir_r()` return a pointer to an object of type `struct dirent`. When an error is encountered, a null pointer is returned and `errno` is set to indicate the error. When the end of the directory is encountered, a null pointer is returned and `errno` is not changed.

The standard-conforming `readdir_r()` returns 0 if the end of the directory is encountered or a directory entry is stored in the structure referenced by *entry*. Otherwise, an error number is returned to indicate the failure.

ERRORS

The `readdir()` and `readdir_r()` functions will fail if:

`EOVERFLOW` One of the values in the structure to be returned cannot be represented correctly.

The `readdir()` and `readdir_r()` functions may fail if:

`EBADF` The *dirp* argument does not refer to an open directory stream.

`ENOENT` The current position of the directory stream is invalid.

USAGE

The `readdir()` and `readdir_r()` functions should be used in conjunction with `opendir()`, `closedir()`, and `rewinddir()` to examine the contents of the directory. Since `readdir()` and the default `readdir_r()` return a null pointer both at the end of the directory and on error, an application wanting to check for error situations should set `errno` to 0 before calling either of these functions. If `errno` is set to non-zero on return, an error occurred.

The standard-conforming `readdir_r()` returns the error number if an error occurred. It returns 0 on success (including reaching the end of the directory stream).

readdir(3C)

The `readdir()` and `readdir_r()` functions have transitional interfaces for 64-bit file offsets. See 1f64(5).

EXAMPLES

EXAMPLE 1 Search the current directory for the entry *name*.

The following sample program will search the current directory for each of the arguments supplied on the command line:

```
#include <sys/types.h>
#include <dirent.h>
#include <errno.h>
#include <stdio.h>
#include <strings.h>

static void lookup(const char *arg)
{
    DIR *dirp;
    struct dirent *dp;

    if ((dirp = opendir(".")) == NULL) {
        perror("couldn't open '.'");
        return;
    }

    do {
        errno = 0;
        if ((dp = readdir(dirp)) != NULL) {
            if (strcmp(dp->d_name, arg) != 0)
                continue;

            (void) printf("found %s\n", arg);
            (void) closedir(dirp);
            return;
        }
    } while (dp != NULL);

    if (errno != 0)
        perror("error reading directory");
    else
        (void) printf("failed to find %s\n", arg);
    (void) closedir(dirp);
    return;
}

int main(int argc, char *argv[])
{
    int i;
    for (i = 1; i < argc; i++)
        lookup(argv[i]);
    return (0);
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	<code>readdir()</code> is Unsafe; <code>readdir_r()</code> is Safe.

SEE ALSO `fork(2)`, `lstat(2)`, `symlink(2)`, `Intro(3)`, `closedir(3C)`, `opendir(3C)`, `rewinddir(3C)`, `seekdir(3C)`, `attributes(5)`, `lf64(5)`, `standards(5)`

NOTES When compiling multithreaded programs, see `Intro(3)`, *Notes On Multithreaded Applications*.

Solaris 2.4 and earlier releases provided a `readdir_r()` interface 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 standard-conforming interface.

For POSIX.1c-conforming applications, the `_POSIX_PTHREAD_SEMANTICS` and `_REENTRANT` flags are automatically turned on by defining the `_POSIX_C_SOURCE` flag with a value $\geq 199506L$.

readdir(3UCB)

NAME	readdir – read a directory entry																				
SYNOPSIS	<pre><code>/usr/ucb/cc[flag ...] file ... #include <sys/types.h> #include <sys/dir.h> struct direct *readdir (dirp); DIR *dirp;</code></pre>																				
DESCRIPTION	<p>The <code>readdir()</code> function returns a pointer to a structure representing the directory entry at the current position in the directory stream to which <code>dirp</code> refers, and positions the directory stream at the next entry, except on read-only file systems. It returns a NULL pointer upon reaching the end of the directory stream, or upon detecting an invalid location in the directory. The <code>readdir()</code> function shall not return directory entries containing empty names. It is unspecified whether entries are returned for dot (<code>.</code>) or dot-dot (<code>..</code>). The pointer returned by <code>readdir()</code> points to data that may be overwritten by another call to <code>readdir()</code> on the same directory stream. This data shall not be overwritten by another call to <code>readdir()</code> on a different directory stream. The <code>readdir()</code> function may buffer several directory entries per actual read operation. The <code>readdir()</code> function marks for update the <code>st_atime</code> field of the directory each time the directory is actually read.</p>																				
RETURN VALUES	The <code>readdir()</code> function returns NULL on failure and sets <code>errno</code> to indicate the error.																				
ERRORS	<p>The <code>readdir()</code> function will fail if one or more of the following are true:</p> <table><tr><td>EAGAIN</td><td>Mandatory file/record locking was set, O_NDELAY or O_NONBLOCK was set, and there was a blocking record lock.</td></tr><tr><td>EAGAIN</td><td>Total amount of system memory available when reading using raw I/O is temporarily insufficient.</td></tr><tr><td>EAGAIN</td><td>No data is waiting to be read on a file associated with a tty device and O_NONBLOCK was set.</td></tr><tr><td>EAGAIN</td><td>No message is waiting to be read on a stream and O_NDELAY or O_NONBLOCK was set.</td></tr><tr><td>EBADF</td><td>The file descriptor determined by the DIR stream is no longer valid. This results if the DIR stream has been closed.</td></tr><tr><td>EBADMSG</td><td>Message waiting to be read on a stream is not a data message.</td></tr><tr><td>EDEADLK</td><td>The read() was going to go to sleep and cause a deadlock to occur.</td></tr><tr><td>EFAULT</td><td><i>buf</i> points to an illegal address.</td></tr><tr><td>EINTR</td><td>A signal was caught during the read() or readv() function.</td></tr><tr><td>EINVAL</td><td>Attempted to read from a stream linked to a multiplexor.</td></tr></table>	EAGAIN	Mandatory file/record locking was set, O_NDELAY or O_NONBLOCK was set, and there was a blocking record lock.	EAGAIN	Total amount of system memory available when reading using raw I/O is temporarily insufficient.	EAGAIN	No data is waiting to be read on a file associated with a tty device and O_NONBLOCK was set.	EAGAIN	No message is waiting to be read on a stream and O_NDELAY or O_NONBLOCK was set.	EBADF	The file descriptor determined by the DIR stream is no longer valid. This results if the DIR stream has been closed.	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	<i>buf</i> points to an illegal address.	EINTR	A signal was caught during the read() or readv() function.	EINVAL	Attempted to read from a stream linked to a multiplexor.
EAGAIN	Mandatory file/record locking was set, O_NDELAY or O_NONBLOCK was set, and there was a blocking record lock.																				
EAGAIN	Total amount of system memory available when reading using raw I/O is temporarily insufficient.																				
EAGAIN	No data is waiting to be read on a file associated with a tty device and O_NONBLOCK was set.																				
EAGAIN	No message is waiting to be read on a stream and O_NDELAY or O_NONBLOCK was set.																				
EBADF	The file descriptor determined by the DIR stream is no longer valid. This results if the DIR stream has been closed.																				
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	<i>buf</i> points to an illegal address.																				
EINTR	A signal was caught during the read() or readv() function.																				
EINVAL	Attempted to read from a stream linked to a multiplexor.																				

readdir(3UCB)

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 SIGTTIN signal or the process group of the process is orphaned.
ENOENT	The current file pointer for the directory is not located at a valid entry.
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	<i>fildev</i> 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.
EOVERFLOW	The value of the <code>direct</code> structure member <code>d_ino</code> cannot be represented in an <code>ino_t</code> .

USAGE The `readdir()` function has a transitional interface for 64-bit file offsets. See [1f64\(5\)](#).

SEE ALSO [getdents\(2\)](#), [readdir\(3C\)](#), [scandir\(3UCB\)](#), [1f64\(5\)](#)

NOTES Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.

realpath(3C)

NAME	realpath – resolve pathname																				
SYNOPSIS	<pre>#include <stdlib.h> char *realpath(const char *restrict file_name, char *restrict resolved_name);</pre>																				
DESCRIPTION	The <code>realpath()</code> function derives, from the pathname pointed to by <code>file_name</code> , an absolute pathname that names the same file, whose resolution does not involve ".", "..", or symbolic links. The generated pathname is stored as a null-terminated string, up to a maximum of <code>{PATH_MAX}</code> bytes (defined in <code>limits.h(3HEAD)</code>), in the buffer pointed to by <code>resolved_name</code> .																				
RETURN VALUES	On successful completion, <code>realpath()</code> returns a pointer to the resolved name. Otherwise, <code>realpath()</code> returns a null pointer and sets <code>errno</code> to indicate the error, and the contents of the buffer pointed to by <code>resolved_name</code> are left in an indeterminate state.																				
ERRORS	<p>The <code>realpath()</code> function will fail if:</p> <table><tr><td>EACCES</td><td>Read or search permission was denied for a component of <code>file_name</code>.</td></tr><tr><td>EINVAL</td><td>Either the <code>file_name</code> or <code>resolved_name</code> argument is a null pointer.</td></tr><tr><td>EIO</td><td>An error occurred while reading from the file system.</td></tr><tr><td>ELOOP</td><td>Too many symbolic links were encountered in resolving <code>file_name</code>.</td></tr><tr><td>ELOOP</td><td>A loop exists in symbolic links encountered during resolution of the <code>file_name</code> argument.</td></tr><tr><td>ENAMETOOLONG</td><td>The <code>file_name</code> argument is longer than <code>{PATH_MAX}</code> or a pathname component is longer than <code>{NAME_MAX}</code>.</td></tr><tr><td>ENOENT</td><td>A component of <code>file_name</code> does not name an existing file or <code>file_name</code> points to an empty string.</td></tr><tr><td>ENOTDIR</td><td>A component of the path prefix is not a directory.</td></tr></table> <p>The <code>realpath()</code> function may fail if:</p> <table><tr><td>ENAMETOOLONG</td><td>Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>{PATH_MAX}</code>.</td></tr><tr><td>ENOMEM</td><td>Insufficient storage space is available.</td></tr></table>	EACCES	Read or search permission was denied for a component of <code>file_name</code> .	EINVAL	Either the <code>file_name</code> or <code>resolved_name</code> argument is a null pointer.	EIO	An error occurred while reading from the file system.	ELOOP	Too many symbolic links were encountered in resolving <code>file_name</code> .	ELOOP	A loop exists in symbolic links encountered during resolution of the <code>file_name</code> argument.	ENAMETOOLONG	The <code>file_name</code> argument is longer than <code>{PATH_MAX}</code> or a pathname component is longer than <code>{NAME_MAX}</code> .	ENOENT	A component of <code>file_name</code> does not name an existing file or <code>file_name</code> points to an empty string.	ENOTDIR	A component of the path prefix is not a directory.	ENAMETOOLONG	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>{PATH_MAX}</code> .	ENOMEM	Insufficient storage space is available.
EACCES	Read or search permission was denied for a component of <code>file_name</code> .																				
EINVAL	Either the <code>file_name</code> or <code>resolved_name</code> argument is a null pointer.																				
EIO	An error occurred while reading from the file system.																				
ELOOP	Too many symbolic links were encountered in resolving <code>file_name</code> .																				
ELOOP	A loop exists in symbolic links encountered during resolution of the <code>file_name</code> argument.																				
ENAMETOOLONG	The <code>file_name</code> argument is longer than <code>{PATH_MAX}</code> or a pathname component is longer than <code>{NAME_MAX}</code> .																				
ENOENT	A component of <code>file_name</code> does not name an existing file or <code>file_name</code> points to an empty string.																				
ENOTDIR	A component of the path prefix is not a directory.																				
ENAMETOOLONG	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <code>{PATH_MAX}</code> .																				
ENOMEM	Insufficient storage space is available.																				
USAGE	<p>The <code>realpath()</code> function operates on null-terminated strings.</p> <p>Execute permission is required for all the directories in the given and the resolved path.</p> <p>The <code>realpath()</code> function might fail to return to the current directory if an error occurs.</p>																				

realpath(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `getcwd(3C)`, `limits.h(3HEAD)`, `sysconf(3C)`, `attributes(5)`, `standards(5)`

reboot(3C)

NAME	reboot – reboot system or halt processor								
SYNOPSIS	<pre>#include <sys/reboot.h> int reboot(int <i>howto</i>, char *<i>bootargs</i>);</pre>								
DESCRIPTION	<p>The <code>reboot()</code> function reboots the system. The <i>howto</i> argument specifies the behavior of the system while rebooting and is a mask constructed by a bitwise-inclusive-OR of flags from the following list:</p> <table><tr><td>RE_AUTOBOOT</td><td>The machine is rebooted from the root filesystem on the default boot device. This is the default behavior. See <code>boot(1M)</code> and <code>kernel(1M)</code>.</td></tr><tr><td>RB_HALT</td><td>The processor is simply halted; no reboot takes place. This option should be used with caution.</td></tr><tr><td>RB_ASKNAME</td><td>Interpreted by the bootstrap program and kernel, causing the user to be asked for pathnames during the bootstrap.</td></tr><tr><td>RB_DUMP</td><td>The system is forced to panic immediately without any further processing and a crash dump is written to the dump device (see <code>dumpadm(1M)</code>) before rebooting.</td></tr></table> <p>Any other <i>howto</i> argument causes the kernel file to boot.</p> <p>The interpretation of the <i>bootargs</i> argument is platform-dependent.</p>	RE_AUTOBOOT	The machine is rebooted from the root filesystem on the default boot device. This is the default behavior. See <code>boot(1M)</code> and <code>kernel(1M)</code> .	RB_HALT	The processor is simply halted; no reboot takes place. This option should be used with caution.	RB_ASKNAME	Interpreted by the bootstrap program and kernel, causing the user to be asked for pathnames during the bootstrap.	RB_DUMP	The system is forced to panic immediately without any further processing and a crash dump is written to the dump device (see <code>dumpadm(1M)</code>) before rebooting.
RE_AUTOBOOT	The machine is rebooted from the root filesystem on the default boot device. This is the default behavior. See <code>boot(1M)</code> and <code>kernel(1M)</code> .								
RB_HALT	The processor is simply halted; no reboot takes place. This option should be used with caution.								
RB_ASKNAME	Interpreted by the bootstrap program and kernel, causing the user to be asked for pathnames during the bootstrap.								
RB_DUMP	The system is forced to panic immediately without any further processing and a crash dump is written to the dump device (see <code>dumpadm(1M)</code>) before rebooting.								
RETURN VALUES	Upon successful completion, <code>reboot()</code> never returns. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.								
ERRORS	The <code>reboot()</code> function will fail if: <table><tr><td>EPERM</td><td>The <code>{PRIV_SYS_CONFIG}</code> privilege is not asserted in the effective set of the calling process.</td></tr></table>	EPERM	The <code>{PRIV_SYS_CONFIG}</code> privilege is not asserted in the effective set of the calling process.						
EPERM	The <code>{PRIV_SYS_CONFIG}</code> privilege is not asserted in the effective set of the calling process.								
SEE ALSO	<code>intro(1M)</code> , <code>boot(1M)</code> , <code>dumpadm(1M)</code> , <code>halt(1M)</code> , <code>init(1M)</code> , <code>kernel(1M)</code> , <code>reboot(1M)</code> , <code>uadmin(2)</code>								

NAME	re_comp, re_exec – compile and execute regular expressions
SYNOPSIS	<pre>#include <re_comp.h> char *re_comp(const char *string); int re_exec(const char *string);</pre>
DESCRIPTION	<p>The re_comp() function converts a regular expression string (RE) into an internal form suitable for pattern matching. The re_exec() function compares the string pointed to by the <i>string</i> argument with the last regular expression passed to re_comp().</p> <p>If re_comp() is called with a null pointer argument, the current regular expression remains unchanged.</p> <p>Strings passed to both re_comp() and re_exec() must be terminated by a null byte, and may include NEWLINE characters.</p> <p>The re_comp() and re_exec() functions support <i>simple regular expressions</i>, which are defined on the regex(5) manual page. The regular expressions of the form <code>\{m\}</code>, <code>\{m,\}</code>, or <code>\{m,n\}</code> are not supported.</p>
RETURN VALUES	<p>The re_comp() function returns a null pointer when the string pointed to by the <i>string</i> argument is successfully converted. Otherwise, a pointer to one of the following error message strings is returned:</p> <pre>No previous regular expression Regular expression too long unmatched \ (missing] too many \ (\) pairs unmatched \)</pre> <p>Upon successful completion, re_exec() returns 1 if <i>string</i> matches the last compiled regular expression. Otherwise, re_exec() returns 0 if <i>string</i> fails to match the last compiled regular expression, and -1 if the compiled regular expression is invalid (indicating an internal error).</p>
ERRORS	No errors are defined.
USAGE	For portability to implementations conforming to X/Open standards prior to SUS, regcomp(3C) and regex(3C) are preferred to these functions. See standards(5) .
SEE ALSO	grep(1), regcmp(1), regcmp(3C) , regcomp(3C) , regex(3C) , regex(3GEN) , regex(5) , standards(5)

regcmp(3C)

NAME	regcmp, regex – compile and execute regular expression												
SYNOPSIS	<pre>#include <libgen.h> char *regcmp(const char *string1, /* char *string2 */ ..., int /* (char*)0 */); char *regex(const char *re, const char *subject, /* char *ret0 */ ...); extern char *__loc1;</pre>												
DESCRIPTION	<p>The <code>regcmp()</code> function compiles a regular expression (consisting of the concatenated arguments) and returns a pointer to the compiled form. The <code>malloc(3C)</code> function is used to create space for the compiled form. It is the user's responsibility to free unneeded space so allocated. A <code>NULL</code> return from <code>regcmp()</code> indicates an incorrect argument. <code>regcmp(1)</code> has been written to generally preclude the need for this routine at execution time.</p> <p>The <code>regex()</code> function executes a compiled pattern against the subject string. Additional arguments are passed to receive values back. The <code>regex()</code> function returns <code>NULL</code> on failure or a pointer to the next unmatched character on success. A global character pointer <code>__loc1</code> points to where the match began. The <code>regcmp()</code> and <code>regex()</code> functions were mostly borrowed from the editor <code>ed(1)</code>; however, the syntax and semantics have been changed slightly. The following are the valid symbols and associated meanings.</p> <table><tr><td>[] * . ^</td><td>This group of symbols retains its meaning as described on the <code>regex(5)</code> manual page.</td></tr><tr><td>\$</td><td>Matches the end of the string; <code>\n</code> matches a newline.</td></tr><tr><td>-</td><td>Within brackets the minus means <i>through</i>. For example, <code>[a-z]</code> is equivalent to <code>[abcd . . .xyz]</code>. The <code>-</code> can appear as itself only if used as the first or last character. For example, the character class expression <code>[]-</code> matches the characters <code>]</code> and <code>-</code>.</td></tr><tr><td>+</td><td>A regular expression followed by <code>+</code> means <i>one or more times</i>. For example, <code>[0-9]+</code> is equivalent to <code>[0-9][0-9]*</code>.</td></tr><tr><td>{m} {m,} {m,u}</td><td>Integer values enclosed in <code>{ }</code> indicate the number of times the preceding regular expression is to be applied. The value <code>m</code> is the minimum number and <code>u</code> is a number, less than 256, which is the maximum. If only <code>m</code> is present (that is, <code>{m}</code>), it indicates the exact number of times the regular expression is to be applied. The value <code>{m,}</code> is analogous to <code>{m,infinity}</code>. The plus (<code>+</code>) and star (<code>*</code>) operations are equivalent to <code>{1,}</code> and <code>{0,}</code> respectively.</td></tr><tr><td>(...)\$n</td><td>The value of the enclosed regular expression is to be returned. The value will be stored in the $(n+1)$th argument following the subject argument. At most, ten enclosed regular expressions are allowed. The <code>regex()</code> function makes its assignments unconditionally.</td></tr></table>	[] * . ^	This group of symbols retains its meaning as described on the <code>regex(5)</code> manual page.	\$	Matches the end of the string; <code>\n</code> matches a newline.	-	Within brackets the minus means <i>through</i> . For example, <code>[a-z]</code> is equivalent to <code>[abcd . . .xyz]</code> . The <code>-</code> can appear as itself only if used as the first or last character. For example, the character class expression <code>[]-</code> matches the characters <code>]</code> and <code>-</code> .	+	A regular expression followed by <code>+</code> means <i>one or more times</i> . For example, <code>[0-9]+</code> is equivalent to <code>[0-9][0-9]*</code> .	{m} {m,} {m,u}	Integer values enclosed in <code>{ }</code> indicate the number of times the preceding regular expression is to be applied. The value <code>m</code> is the minimum number and <code>u</code> is a number, less than 256, which is the maximum. If only <code>m</code> is present (that is, <code>{m}</code>), it indicates the exact number of times the regular expression is to be applied. The value <code>{m,}</code> is analogous to <code>{m,infinity}</code> . The plus (<code>+</code>) and star (<code>*</code>) operations are equivalent to <code>{1,}</code> and <code>{0,}</code> respectively.	(...)\$n	The value of the enclosed regular expression is to be returned. The value will be stored in the $(n+1)$ th argument following the subject argument. At most, ten enclosed regular expressions are allowed. The <code>regex()</code> function makes its assignments unconditionally.
[] * . ^	This group of symbols retains its meaning as described on the <code>regex(5)</code> manual page.												
\$	Matches the end of the string; <code>\n</code> matches a newline.												
-	Within brackets the minus means <i>through</i> . For example, <code>[a-z]</code> is equivalent to <code>[abcd . . .xyz]</code> . The <code>-</code> can appear as itself only if used as the first or last character. For example, the character class expression <code>[]-</code> matches the characters <code>]</code> and <code>-</code> .												
+	A regular expression followed by <code>+</code> means <i>one or more times</i> . For example, <code>[0-9]+</code> is equivalent to <code>[0-9][0-9]*</code> .												
{m} {m,} {m,u}	Integer values enclosed in <code>{ }</code> indicate the number of times the preceding regular expression is to be applied. The value <code>m</code> is the minimum number and <code>u</code> is a number, less than 256, which is the maximum. If only <code>m</code> is present (that is, <code>{m}</code>), it indicates the exact number of times the regular expression is to be applied. The value <code>{m,}</code> is analogous to <code>{m,infinity}</code> . The plus (<code>+</code>) and star (<code>*</code>) operations are equivalent to <code>{1,}</code> and <code>{0,}</code> respectively.												
(...)\$n	The value of the enclosed regular expression is to be returned. The value will be stored in the $(n+1)$ th argument following the subject argument. At most, ten enclosed regular expressions are allowed. The <code>regex()</code> function makes its assignments unconditionally.												

(...) Parentheses are used for grouping. An operator, for example, *, +, { }, can work on a single character or a regular expression enclosed in parentheses. For example, (a*(cb+)*)\$0. By necessity, all the above defined symbols are special. They must, therefore, be escaped with a \ (backslash) to be used as themselves.

EXAMPLES

EXAMPLE 1 Example matching a leading newline in the subject string.

The following example matches a leading newline in the subject string pointed at by cursor.

```
char *cursor, *newcursor, *ptr;
. . .
newcursor = regex((ptr = regcmp("\n", (char *)0)), cursor);
free(ptr);
```

The following example matches through the string `Testing3` and returns the address of the character after the last matched character (the "4"). The string `Testing3` is copied to the character array `ret0`.

```
char ret0[9];
char *newcursor, *name;
. . .
name = regcmp("[A-Za-z][A-Za-z0-9]{0,7}$0", (char *)0);
newcursor = regex(name, "012Testing345", ret0);
```

The following example applies a precompiled regular expression in `file.i` (see `regcmp(1)`) against `string`.

```
#include "file.i"
char *string, *newcursor;
. . .
newcursor = regex(name, string);
```

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO

`ed(1)`, `regcmp(1)`, `malloc(3C)`, `attributes(5)`, `regexp(5)`

NOTES

The user program may run out of memory if `regcmp()` is called iteratively without freeing the vectors no longer required.

When compiling multithreaded applications, the `_REENTRANT` flag must be defined on the compile line. This flag should only be used in multithreaded applications.

regcomp(3C)

NAME	regcomp, regexec, regerror, regfree – regular expression matching								
SYNOPSIS	<pre>#include <sys/types.h> #include <regex.h> int regcomp(regex_t *restrict preg, const char *restrict pattern, int cflags); int regexec(const regex_t *restrict preg, const char *restrict string, size_t nmatch, regmatch_t pmatch[restrict], int eflags); size_t regerror(int errcode, const regex_t *restrict preg, char *restrict errbuf, size_t errbuf_size); void regfree(regex_t *preg);</pre>								
DESCRIPTION	<p>These functions interpret <i>basic</i> and <i>extended</i> regular expressions (described on the regex(5) manual page).</p> <p>The structure type <code>regex_t</code> contains at least the following member:</p> <table><tr><td><code>size_t re_nsub</code></td><td>Number of parenthesised subexpressions.</td></tr></table> <p>The structure type <code>regmatch_t</code> contains at least the following members:</p> <table><tr><td><code>regoff_t rm_so</code></td><td>Byte offset from start of <i>string</i> to start of substring.</td></tr><tr><td><code>regoff_t rm_eo</code></td><td>Byte offset from start of <i>string</i> of the first character after the end of substring.</td></tr></table>	<code>size_t re_nsub</code>	Number of parenthesised subexpressions.	<code>regoff_t rm_so</code>	Byte offset from start of <i>string</i> to start of substring.	<code>regoff_t rm_eo</code>	Byte offset from start of <i>string</i> of the first character after the end of substring.		
<code>size_t re_nsub</code>	Number of parenthesised subexpressions.								
<code>regoff_t rm_so</code>	Byte offset from start of <i>string</i> to start of substring.								
<code>regoff_t rm_eo</code>	Byte offset from start of <i>string</i> of the first character after the end of substring.								
regcomp()	<p>The <code>regcomp()</code> function will compile the regular expression contained in the string pointed to by the <i>pattern</i> argument and place the results in the structure pointed to by <i>preg</i>. The <i>cflags</i> argument is the bitwise inclusive OR of zero or more of the following flags, which are defined in the header <code><regex.h></code>:</p> <table><tr><td><code>REG_EXTENDED</code></td><td>Use Extended Regular Expressions.</td></tr><tr><td><code>REG_ICASE</code></td><td>Ignore case in match.</td></tr><tr><td><code>REG_NOSUB</code></td><td>Report only success/fail in <code>regexec()</code>.</td></tr><tr><td><code>REG_NEWLINE</code></td><td>Change the handling of NEWLINE characters, as described in the text.</td></tr></table> <p>The default regular expression type for <i>pattern</i> is a Basic Regular Expression. The application can specify Extended Regular Expressions using the <code>REG_EXTENDED</code> <i>cflags</i> flag.</p> <p>If the <code>REG_NOSUB</code> flag was not set in <i>cflags</i>, then <code>regcomp()</code> will set <code>re_nsub</code> to the number of parenthesised subexpressions (delimited by <code>\(\)</code> in basic regular expressions or <code>()</code> in extended regular expressions) found in <i>pattern</i>.</p>	<code>REG_EXTENDED</code>	Use Extended Regular Expressions.	<code>REG_ICASE</code>	Ignore case in match.	<code>REG_NOSUB</code>	Report only success/fail in <code>regexec()</code> .	<code>REG_NEWLINE</code>	Change the handling of NEWLINE characters, as described in the text.
<code>REG_EXTENDED</code>	Use Extended Regular Expressions.								
<code>REG_ICASE</code>	Ignore case in match.								
<code>REG_NOSUB</code>	Report only success/fail in <code>regexec()</code> .								
<code>REG_NEWLINE</code>	Change the handling of NEWLINE characters, as described in the text.								

regexec() The `regexec()` function compares the null-terminated string specified by *string* with the compiled regular expression *preg* initialized by a previous call to `regcomp()`. The *eflags* argument is the bitwise inclusive OR of zero or more of the following flags, which are defined in the header `<regex.h>`:

REG_NOTBOL The first character of the string pointed to by *string* is not the beginning of the line. Therefore, the circumflex character (^), when taken as a special character, will not match the beginning of *string*.

REG_NOTEOL The last character of the string pointed to by *string* is not the end of the line. Therefore, the dollar sign (\$), when taken as a special character, will not match the end of *string*.

If *nmatch* is zero or **REG_NOSUB** was set in the *cflags* argument to `regcomp()`, then `regexec()` will ignore the *pmatch* argument. Otherwise, the *pmatch* argument must point to an array with at least *nmatch* elements, and `regexec()` will fill in the elements of that array with offsets of the substrings of *string* that correspond to the parenthesised subexpressions of *pattern*: *pmatch*[*i*].*rm_so* will be the byte offset of the beginning and *pmatch*[*i*].*rm_eo* will be one greater than the byte offset of the end of substring *i*. (Subexpression *i* begins at the *i*th matched open parenthesis, counting from 1.) Offsets in *pmatch*[0] identify the substring that corresponds to the entire regular expression. Unused elements of *pmatch* up to *pmatch*[*nmatch*-1] will be filled with -1. If there are more than *nmatch* subexpressions in *pattern* (*pattern* itself counts as a subexpression), then `regexec()` will still do the match, but will record only the first *nmatch* substrings.

When matching a basic or extended regular expression, any given parenthesised subexpression of *pattern* might participate in the match of several different substrings of *string*, or it might not match any substring even though the pattern as a whole did match. The following rules are used to determine which substrings to report in *pmatch* when matching regular expressions:

1. If subexpression *i* in a regular expression is not contained within another subexpression, and it participated in the match several times, then the byte offsets in *pmatch*[*i*] will delimit the last such match.
2. If subexpression *i* is not contained within another subexpression, and it did not participate in an otherwise successful match, the byte offsets in *pmatch*[*i*] will be -1. A subexpression does not participate in the match when:

* or `\{\}` appears immediately after the subexpression in a basic regular expression, or *, ?, or `{ }` appears immediately after the subexpression in an extended regular expression, and the subexpression did not match (matched zero times)

or

| is used in an extended regular expression to select this subexpression or another, and the other subexpression matched.

regcomp(3C)

3. If subexpression *i* is contained within another subexpression *j*, and *i* is not contained within any other subexpression that is contained within *j*, and a match of subexpression *j* is reported in *pmatch* [*j*], then the match or non-match of subexpression *i* reported in *pmatch* [*i*] will be as described in 1. and 2. above, but within the substring reported in *pmatch* [*j*] rather than the whole string.
4. If subexpression *i* is contained in subexpression *j*, and the byte offsets in *pmatch* [*j*] are -1 , then the pointers in *pmatch* [*i*] also will be -1 .
5. If subexpression *i* matched a zero-length string, then both byte offsets in *pmatch* [*i*] will be the byte offset of the character or NULL terminator immediately following the zero-length string.

If, when `regexec()` is called, the locale is different from when the regular expression was compiled, the result is undefined.

If `REG_NEWLINE` is not set in *flags*, then a NEWLINE character in *pattern* or *string* will be treated as an ordinary character. If `REG_NEWLINE` is set, then newline will be treated as an ordinary character except as follows:

1. A NEWLINE character in *string* will not be matched by a period outside a bracket expression or by any form of a non-matching list.
2. A circumflex (^) in *pattern*, when used to specify expression anchoring will match the zero-length string immediately after a newline in *string*, regardless of the setting of `REG_NOTBOL`.
3. A dollar-sign (\$) in *pattern*, when used to specify expression anchoring, will match the zero-length string immediately before a newline in *string*, regardless of the setting of `REG_NOTEOL`.

regfree() The `regfree()` function frees any memory allocated by `regcomp()` associated with *preg*.

The following constants are defined as error return values:

<code>REG_NOMATCH</code>	The <code>regexec()</code> function failed to match.
<code>REG_BADPAT</code>	Invalid regular expression.
<code>REG_ECOLLATE</code>	Invalid collating element referenced.
<code>REG_ECTYPE</code>	Invalid character class type referenced.
<code>REG_EESCAPE</code>	Trailing <code>\</code> in pattern.
<code>REG_ESUBREG</code>	Number in <code>\digit</code> invalid or in error.
<code>REG_EBRACK</code>	[] imbalance.
<code>REG_ENOSYS</code>	The function is not supported.
<code>REG_EPAREN</code>	<code>\(\)</code> or <code>()</code> imbalance.

REG_EBRACE	\{ \} imbalance.
REG_BADBR	Content of \{ \} invalid: not a number, number too large, more than two numbers, first larger than second.
REG_ERANGE	Invalid endpoint in range expression.
REG_ESPACE	Out of memory.
REG_BADRPT	?, * or + not preceded by valid regular expression.

regerror() The `regerror()` function provides a mapping from error codes returned by `regcomp()` and `regexexec()` to unspecified printable strings. It generates a string corresponding to the value of the *errcode* argument, which must be the last non-zero value returned by `regcomp()` or `regexexec()` with the given value of *preg*. If *errcode* is not such a value, an error message indicating that the error code is invalid is returned.

If *preg* is a `NULL` pointer, but *errcode* is a value returned by a previous call to `regexexec()` or `regcomp()`, the `regerror()` still generates an error string corresponding to the value of *errcode*.

If the *errbuf_size* argument is not zero, `regerror()` will place the generated string into the buffer of size *errbuf_size* bytes pointed to by *errbuf*. If the string (including the terminating `NULL`) cannot fit in the buffer, `regerror()` will truncate the string and null-terminate the result.

If *errbuf_size* is zero, `regerror()` ignores the *errbuf* argument, and returns the size of the buffer needed to hold the generated string.

If the *preg* argument to `regexexec()` or `regfree()` is not a compiled regular expression returned by `regcomp()`, the result is undefined. A *preg* is no longer treated as a compiled regular expression after it is given to `regfree()`.

See `regex(5)` for BRE (Basic Regular Expression) Anchoring.

RETURN VALUES On successful completion, the `regcomp()` function returns 0. Otherwise, it returns an integer value indicating an error as described in `<regex.h>`, and the content of *preg* is undefined.

On successful completion, the `regexexec()` function returns 0. Otherwise it returns `REG_NOMATCH` to indicate no match, or `REG_ENOSYS` to indicate that the function is not supported.

Upon successful completion, the `regerror()` function returns the number of bytes needed to hold the entire generated string. Otherwise, it returns 0 to indicate that the function is not implemented.

The `regfree()` function returns no value.

ERRORS No errors are defined.

regcomp(3C)

USAGE An application could use:

```
regerror (code, preg, (char *)NULL, (size_t)0)
```

to find out how big a buffer is needed for the generated string, `malloc` a buffer to hold the string, and then call `regerror()` again to get the string (see [malloc\(3C\)](#)). Alternately, it could allocate a fixed, static buffer that is big enough to hold most strings, and then use `malloc()` to allocate a larger buffer if it finds that this is too small.

EXAMPLES **EXAMPLE 1** Example to match string against the extended regular expression in pattern.

```
#include <regex.h>
/*
 * Match string against the extended regular expression in
 * pattern, treating errors as no match.
 *
 * return 1 for match, 0 for no match
 */

int
match(const char *string, char *pattern)
{
    int status;
    regex_t re;
    if (regcomp(&re, pattern, REG_EXTENDED | REG_NOSUB) != 0) {
        return(0);    /* report error */
    }
    status = regexec(&re, string, (size_t) 0, NULL, 0);
    regfree(&re);
    if (status != 0) {
        return(0);    /* report error */
    }
    return(1);
}
```

The following demonstrates how the `REG_NOTBOL` flag could be used with `regexec()` to find all substrings in a line that match a pattern supplied by a user. (For simplicity of the example, very little error checking is done.)

```
(void) regcomp (&re, pattern, 0);
/* this call to regexec() finds the first match on the line */
error = regexec (&re, &buffer[0], 1, &pm, 0);
while (error == 0) {    /* while matches found */
    /* substring found between pm.rm_so and pm.rm_eo */
    /* This call to regexec() finds the next match */
    error = regexec (&re, buffer + pm.rm_eo, 1, &pm, REG_NOTBOL);
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

SEE ALSO `fnmatch(3C)`, `glob(3C)`, `malloc(3C)`, `setlocale(3C)`, `attributes(5)`, `standards(5)`, `regex(5)`

NOTES The `regcomp()` function can be used safely in a multithreaded application as long as `setlocale(3C)` is not being called to change the locale.

remove(3C)

NAME remove – remove file

SYNOPSIS #include <stdio.h>
int **remove**(const char *path);

DESCRIPTION The `remove()` function causes the file or empty directory whose name is the string pointed to by *path* to be no longer accessible by that name. A subsequent attempt to open that file using that name will fail, unless the file is created anew.

For files, `remove()` is identical to `unlink()`. For directories, `remove()` is identical to `rmdir()`.

See `rmdir(2)` and `unlink(2)` for a detailed list of failure conditions.

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

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `rmdir(2)`, `unlink(2)`, `attributes(5)`, `standards(5)`

NAME	rewind – reset file position indicator in a stream						
SYNOPSIS	<pre>#include <stdio.h> void rewind(FILE *<i>stream</i>);</pre>						
DESCRIPTION	<p>The call:</p> <pre>rewind(<i>stream</i>)</pre> <p>is equivalent to:</p> <pre>(void) fseek(<i>stream</i>, 0L, SEEK_SET)</pre> <p>except that <code>rewind()</code> also clears the error indicator.</p>						
RETURN VALUES	The <code>rewind()</code> function returns no value.						
ERRORS	Refer to fseek(3C) with the exception of <code>EINVAL</code> which does not apply.						
USAGE	Because <code>rewind()</code> does not return a value, an application wishing to detect errors should clear <code>errno</code> , then call <code>rewind()</code> , and if <code>errno</code> is non-zero, assume an error has occurred.						
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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	fseek(3C) , attributes(5) , standards(5)						

rewinddir(3C)

NAME	rewinddir – reset position of directory stream to the beginning of a directory						
SYNOPSIS	<pre>#include <sys/types.h> #include <dirent.h> void rewinddir(DIR *dirp);</pre>						
DESCRIPTION	<p>The <code>rewinddir()</code> function resets the position of the directory stream to which <i>dirp</i> refers to the beginning of the directory. It also causes the directory stream to refer to the current state of the corresponding directory, as a call to <code>opendir(3C)</code> would have done. If <i>dirp</i> does not refer to a directory stream, the effect is undefined.</p> <p>After a call to the <code>fork(2)</code> function, either the parent or child (but not both) may continue processing the directory stream using <code>readdir(3C)</code>, <code>rewinddir()</code> or <code>seekdir(3C)</code>. If both the parent and child processes use these functions, the result is undefined.</p>						
RETURN VALUES	The <code>rewinddir()</code> function does not return a value.						
ERRORS	No errors are defined.						
USAGE	The <code>rewinddir()</code> function should be used in conjunction with <code>opendir()</code> , <code>readdir()</code> , and <code>closedir(3C)</code> to examine the contents of the directory. This method is recommended for portability.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	<code>fork(2)</code> , <code>closedir(3C)</code> , <code>opendir(3C)</code> , <code>readdir(3C)</code> , <code>seekdir(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	rwlock, rwlock_init, rwlock_destroy, rw_rdlock, rw_wrlock, rw_tryrdlock, rw_trywrlock, rw_unlock – multiple readers, single writer locks				
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <synch.h> int rwlock_init(rwlock_t *rwlp, int type, void * arg); int rwlock_destroy(rwlock_t *rwlp); int rw_rdlock(rwlock_t *rwlp); int rw_wrlock(rwlock_t *rwlp); int rw_unlock(rwlock_t *rwlp); int rw_tryrdlock(rwlock_t *rwlp); int rw_trywrlock(rwlock_t *rwlp);</pre>				
DESCRIPTION	<p>Many threads can have simultaneous read-only access to data, while only one thread can have write access at any given time. Multiple read access with single write access is controlled by locks, which are generally used to protect data that is frequently searched.</p> <p>Readers/writer locks can synchronize threads in this process and other processes if they are allocated in writable memory and shared among cooperating processes (see <code>mmap(2)</code>), and are initialized for this purpose.</p> <p>Additionally, readers/writer locks must be initialized prior to use. <code>rwlock_init()</code> The readers/writer lock pointed to by <code>rwlp</code> is initialized by <code>rwlock_init()</code>. A readers/writer lock is capable of having several types of behavior, which is specified by <code>type</code>. <code>arg</code> is currently not used, although a future type may define new behavior parameters by way of <code>arg</code>.</p> <p><code>type</code> may be one of the following:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>USYNC_PROCESS</code></td> <td>The readers/writer lock can synchronize threads in this process and other processes. The readers/writer lock should be initialized by only one process. <code>arg</code> is ignored. A readers/writer lock initialized with this type, must be allocated in memory shared between processes, i.e. either in Sys V shared memory (see <code>shmop(2)</code>) or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and to not allocate it in such shared memory.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;"><code>USYNC_THREAD</code></td> <td>The readers/writer lock can synchronize threads in this process, only. <code>arg</code> is ignored.</td> </tr> </table>	<code>USYNC_PROCESS</code>	The readers/writer lock can synchronize threads in this process and other processes. The readers/writer lock should be initialized by only one process. <code>arg</code> is ignored. A readers/writer lock initialized with this type, must be allocated in memory shared between processes, i.e. either in Sys V shared memory (see <code>shmop(2)</code>) or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and to not allocate it in such shared memory.	<code>USYNC_THREAD</code>	The readers/writer lock can synchronize threads in this process, only. <code>arg</code> is ignored.
<code>USYNC_PROCESS</code>	The readers/writer lock can synchronize threads in this process and other processes. The readers/writer lock should be initialized by only one process. <code>arg</code> is ignored. A readers/writer lock initialized with this type, must be allocated in memory shared between processes, i.e. either in Sys V shared memory (see <code>shmop(2)</code>) or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and to not allocate it in such shared memory.				
<code>USYNC_THREAD</code>	The readers/writer lock can synchronize threads in this process, only. <code>arg</code> is ignored.				

rwlock(3C)

Additionally, readers/writer locks can be initialized by allocation in zeroed memory. A type of `USYNC_THREAD` is assumed in this case. Multiple threads must not simultaneously initialize the same readers/writer lock. And a readers/writer lock must not be re-initialized while in use by other threads.

The following are default readers/writer lock initialization (intra-process):

```
rwlock_t rwp;  
rwlock_init(&rwp, NULL, NULL);  
OR  
rwlock_init(&rwp, USYNC_THREAD, NULL);  
OR  
rwlock_t rwp = DEFAULTRWLOCK;
```

The following is a customized readers/writer lock initialization (inter-process):

```
rwlock_init(&rwp, USYNC_PROCESS, NULL);
```

Any state associated with the readers/writer lock pointed to by *rwp* are destroyed by `rwlock_destroy()` and the readers/writer lock storage space is not released.

`rw_rdlock()` gets a read lock on the readers/writer lock pointed to by *rwp*. If the readers/writer lock is currently locked for writing, the calling thread blocks until the write lock is freed. Multiple threads may simultaneously hold a read lock on a readers/writer lock.

`rw_tryrdlock()` tries to get a read lock on the readers/writer lock pointed to by *rwp*. If the readers/writer lock is locked for writing, it returns an error; otherwise, the read lock is acquired.

`rw_wrlock()` gets a write lock on the readers/writer lock pointed to by *rwp*. If the readers/writer lock is currently locked for reading or writing, the calling thread blocks until all the read and write locks are freed. At any given time, only one thread may have a write lock on a readers/writer lock.

`rw_trywrlock()` tries to get a write lock on the readers/writer lock pointed to by *rwp*. If the readers/writer lock is currently locked for reading or writing, it returns an error.

`rw_unlock()` unlocks a readers/writer lock pointed to by *rwp*, if the readers/writer lock is locked and the calling thread holds the lock for either reading or writing. One of the other threads that is waiting for the readers/writer lock to be freed will be unblocked, provided there is other waiting threads. If the calling thread does not hold the lock for either reading or writing, no error status is returned, and the program's behavior is unknown.

RETURN VALUES

If successful, these functions return 0. Otherwise, a non-zero value is returned to indicate the error.

ERRORS

The `rwlock_init()` function will fail if:

EINVAL type is invalid.

The `rw_tryrdlock()` or `rw_trywrlock()` functions will fail if:

EBUSY The reader or writer lock pointed to by *rwlp* was already locked.

These functions may fail if:

EFAULT *rwlp* or *arg* points to an illegal address.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `mmap(2)`, `attributes(5)`

NOTES These interfaces also available by way of:

```
#include <thread.h>
```

If multiple threads are waiting for a readers/writer lock, the acquisition order is random by default. However, some implementations may bias acquisition order to avoid depriving writers. The current implementation favors writers over readers.

scandir(3C)

NAME	scandir, alphasort – scan a directory						
SYNOPSIS	<pre>#include <sys/types.h> #include <dirent.h> int scandir(const char *<i>dirname</i>, struct dirent *(*<i>namelist</i>[]), int (*<i>select</i>)(const struct dirent *), int (*<i>dcomp</i>)(const struct dirent **, const struct dirent **)); int alphasort(const struct dirent **<i>d1</i>, const struct dirent **<i>d2</i>);</pre>						
DESCRIPTION	<p>The <code>scandir()</code> function reads the directory <code>dirname</code> using <code>readdir(3C)</code> and builds an array of pointers to directory entries using <code>malloc(3C)</code>. The <code>namelist</code> argument is a pointer to an array of structure pointers. The <code>select</code> argument is a pointer to a routine that is called with a pointer to a directory entry and returns a non-zero value if the directory entry is included in the array. If this pointer is <code>NULL</code>, then all the directory entries are included. The <code>dcomp</code> argument is a pointer to a routine that is passed to <code>qsort(3C)</code>, which sorts the completed array. If this pointer is <code>NULL</code>, the array is not sorted.</p> <p>The <code>alphasort()</code> function can be used as the <code>dcomp()</code> function parameter for the <code>scandir()</code> function to sort the directory entries into alphabetical order, as if by the <code>strcoll(3C)</code> function. Its arguments are the two directory entries to compare.</p>						
RETURN VALUES	<p>The <code>scandir()</code> function returns the number of entries in the array and a pointer to the array through the <code>namelist</code> argument. When an error is encountered, <code>scandir()</code> returns -1 and <code>errno</code> is set to indicate the error.</p> <p>The <code>alphasort()</code> function returns an integer greater than, equal to, or less than 0 if the directory entry name pointed to by <code>d1</code> is greater than, equal to, or less than the directory entry name pointed to by <code>d2</code> when both are interpreted as appropriate to the current locale. There is no return value reserved to indicate an error.</p>						
ERRORS	<p>The <code>scandir()</code> function will fail if:</p> <table><tr><td><code>E_OVERFLOW</code></td><td>The number of directory entries exceeds the number that can be represented by an <code>int</code>.</td></tr></table>	<code>E_OVERFLOW</code>	The number of directory entries exceeds the number that can be represented by an <code>int</code> .				
<code>E_OVERFLOW</code>	The number of directory entries exceeds the number that can be represented by an <code>int</code> .						
USAGE	The <code>scandir()</code> and <code>alphasort()</code> functions have transitional interfaces for 64-bit file offsets. See <code>lf64(5)</code> .						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Stable</td></tr><tr><td>MT-Level</td><td>See below.</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Stable	MT-Level	See below.
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Stable						
MT-Level	See below.						
	The <code>scandir()</code> function is Unsafe. The <code>alphasort()</code> function is Safe.						

scandir(3C)

SEE ALSO malloc(3C), qsort(3C), readdir(3C), strcoll(3C), attributes(5), lf64(5)

scandir(3UCB)

NAME	scandir, alphasort – scan a directory
SYNOPSIS	<pre>/usr/ucb/cc [flag...] file... #include <sys/types.h> #include <sys/dir.h> int scandir (dirname, namelist, select, dcomp) ; char *dirname; struct direct *(*namelist[]); int (*select(.), (*dcomp)()); int alphasort (d1, d2) ; struct direct **d1, **d2;</pre>
DESCRIPTION	<p>The <code>scandir()</code> function reads the directory <i>dirname</i> and builds an array of pointers to directory entries using <code>malloc(3C)</code>. The second parameter is a pointer to an array of structure pointers. The third parameter is a pointer to a routine which is called with a pointer to a directory entry and should return a non zero value if the directory entry should be included in the array. If this pointer is <code>NULL</code>, then all the directory entries will be included. The last argument is a pointer to a routine which is passed to <code>qsort(3C)</code>, which sorts the completed array. If this pointer is <code>NULL</code>, the array is not sorted.</p> <p>The <code>alphasort()</code> function sorts the array alphabetically.</p>
RETURN VALUES	<p>The <code>scandir()</code> function returns the number of entries in the array and a pointer to the array through the parameter <i>namelist</i>. The <code>scandir()</code> function returns <code>-1</code> if the directory cannot be opened for reading or if <code>malloc(3C)</code> cannot allocate enough memory to hold all the data structures.</p> <p>The <code>alphasort()</code> function returns an integer greater than, equal to, or less than 0 if the directory entry name pointed to by <i>d1</i> is greater than, equal to, or less than the directory entry name pointed to by <i>d2</i>.</p>
USAGE	The <code>scandir()</code> and <code>alphasort()</code> functions have transitional interfaces for 64-bit file offsets. See <code>lf64(5)</code> .
SEE ALSO	<code>getdents(2)</code> , <code>malloc(3C)</code> , <code>qsort(3C)</code> , <code>readdir(3UCB)</code> , <code>readdir(3C)</code> , <code>lf64(5)</code>
NOTES	Use of these functions should be restricted to applications written on BSD platforms. Use of these functions with any of the system libraries or in multithreaded applications is unsupported.

NAME	scanf, fscanf, sscanf, vscanf, vfscanf, vsscanf – convert formatted input
SYNOPSIS	<pre>#include <stdio.h> int scanf(const char *restrict <i>format</i>, ...); int fscanf(FILE *restrict <i>stream</i>, const char *restrict <i>format</i>, ...); int sscanf(const char *restrict <i>s</i>, const char *restrict <i>format</i>, ...); #include <stdarg.h> #include <stdio.h> int vscanf(const char *<i>format</i>, va_list <i>arg</i>); int vfscanf(FILE *<i>stream</i>, const char *<i>format</i>, va_list <i>arg</i>); int vsscanf(const char *<i>s</i>, const char *<i>format</i>, va_list <i>arg</i>);</pre>
DESCRIPTION	<p>The <code>scanf()</code> function reads from the standard input stream <code>stdin</code>.</p> <p>The <code>fscanf()</code> function reads from the named input <i>stream</i>.</p> <p>The <code>sscanf()</code> function reads from the string <i>s</i>.</p> <p>The <code>vscanf()</code>, <code>vfscanf()</code>, and <code>vsscanf()</code> functions are equivalent to the <code>scanf()</code>, <code>fscanf()</code>, and <code>sscanf()</code> functions, respectively, except that instead of being called with a variable number of arguments, they are called with an argument list as defined by the <code><stdarg.h></code> header. These functions do not invoke the <code>va_end()</code> macro. Applications using these functions should call <code>va_end(ap)</code> afterwards to clean up.</p> <p>Each function reads bytes, interprets them according to a <i>format</i>, and stores the results in its arguments. Each expects, as arguments, a control string <i>format</i> described below, and a set of <i>pointer</i> arguments indicating where the converted input should be stored. The result is undefined if there are insufficient arguments for the format. If the format is exhausted while arguments remain, the excess arguments are evaluated but are otherwise ignored.</p> <p>Conversions can be applied to the <i>n</i>th argument after the <i>format</i> in the argument list, rather than to the next unused argument. In this case, the conversion character <code>%</code> (see below) is replaced by the sequence <code>%n\$</code>, where <i>n</i> is a decimal integer in the range [1, <code>NL_ARGMAX</code>]. This feature provides for the definition of format strings that select arguments in an order appropriate to specific languages. In format strings containing the <code>%n\$</code> form of conversion specifications, it is unspecified whether numbered arguments in the argument list can be referenced from the format string more than once.</p> <p>The <i>format</i> can contain either form of a conversion specification, that is, <code>%</code> or <code>%n\$</code>, but the two forms cannot normally be mixed within a single <i>format</i> string. The only exception to this is that <code>%%</code> or <code>%*</code> can be mixed with the <code>%n\$</code> form.</p>

scanf(3C)

The `scanf()` function in all its forms allows for detection of a language-dependent radix character in the input string. The radix character is defined in the program's locale (category `LC_NUMERIC`). In the POSIX locale, or in a locale where the radix character is not defined, the radix character defaults to a period (`.`).

The format is a character string, beginning and ending in its initial shift state, if any, composed of zero or more directives. Each directive is composed of one of the following:

- one or more *white-space characters* (space, tab, newline, vertical-tab or form-feed characters);
- an *ordinary character* (neither `%` nor a white-space character); or
- a *conversion specification*.

Conversion Specifications

Each conversion specification is introduced by the character `%` or the character sequence `%n$`, after which the following appear in sequence:

- An optional assignment-suppressing character `*`.
- An optional non-zero decimal integer that specifies the maximum field width.
- An option length modifier that specifies the size of the receiving object.
- A conversion specifier character that specifies the type of conversion to be applied. The valid conversion characters are described below.

The `scanf()` functions execute each directive of the format in turn. If a directive fails, as detailed below, the function returns. Failures are described as input failures (due to the unavailability of input bytes) or matching failures (due to inappropriate input).

A directive composed of one or more white-space characters is executed by reading input until no more valid input can be read, or up to the first byte which is not a white-space character which remains unread.

A directive that is an ordinary character is executed as follows. The next byte is read from the input and compared with the byte that comprises the directive; if the comparison shows that they are not equivalent, the directive fails, and the differing and subsequent bytes remain unread.

A directive that is a conversion specification defines a set of matching input sequences, as described below for each conversion character. A conversion specification is executed in the following steps:

Input white-space characters (as specified by `isspace(3C)`) are skipped, unless the conversion specification includes a `l`, `c`, `C`, or `n` conversion character.

An item is read from the input unless the conversion specification includes an `n` conversion character. The length of the item read is limited to any specified maximum field width, which is interpreted in either characters or bytes depending on the conversion character. In Solaris default mode, the input item is defined as the longest sequence of input bytes that forms a matching sequence. In some cases, `scanf()` might need to read several extra characters beyond the end of the input item to find

the end of a matching sequence. In C99/SUSv3 mode, the input item is defined as the longest sequence of input bytes that is, or is a prefix of, a matching sequence. With this definition, `scanf()` need only read at most one character beyond the end of the input item. Therefore, in C99/SUSv3 mode, some sequences that are acceptable to `strtod(3C)`, `strtol(3C)`, and similar functions are unacceptable to `scanf()`. In either mode, `scanf()` attempts to push back any excess bytes read using `ungetc(3C)`. Assuming all such attempts succeed, the first byte, if any, after the input item remains unread. If the length of the input item is 0, the conversion fails. This condition is a matching failure unless end-of-file, an encoding error, or a read error prevented input from the stream, in which case it is an input failure.

Except in the case of a % conversion character, the input item (or, in the case of a %*n* conversion specification, the count of input bytes) is converted to a type appropriate to the conversion character. If the input item is not a matching sequence, the execution of the conversion specification fails; this condition is a matching failure. Unless assignment suppression was indicated by a *, the result of the conversion is placed in the object pointed to by the first argument following the *format* argument that has not already received a conversion result if the conversion specification is introduced by %, or in the *n*th argument if introduced by the character sequence %*n*\$. If this object does not have an appropriate type, or if the result of the conversion cannot be represented in the space provided, the behavior is undefined.

Length Modifiers

The length modifiers and their meanings are:

hh	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to signed char or unsigned char.
h	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to short or unsigned short.
l (ell)	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to long or unsigned long; that a following <code>a</code> , <code>A</code> , <code>e</code> , <code>E</code> , <code>f</code> , <code>F</code> , <code>g</code> , or <code>G</code> conversion specifier applies to an argument with type pointer to double; or that a following <code>c</code> , <code>s</code> , or <code>[</code> conversion specifier applies to an argument with type pointer to <code>wchar_t</code> .
ll (ell-ell)	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to long long or unsigned long long.
j	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to <code>intmax_t</code> or <code>uintmax_t</code> .
z	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to <code>size_t</code> or the corresponding signed integer type.
t	Specifies that a following <code>d</code> , <code>i</code> , <code>o</code> , <code>u</code> , <code>x</code> , <code>X</code> , or <code>n</code> conversion specifier applies to an argument with type pointer to <code>ptrdiff_t</code> or the corresponding unsigned type.

scanf(3C)

Conversion Characters

- L Specifies that a following *a*, *A*, *e*, *E*, *f*, *F*, *g*, or *G* conversion specifier applies to an argument with type pointer to long double.
- If a length modifier appears with any conversion specifier other than as specified above, the behavior is undefined.
- The following conversion characters are valid:
- d Matches an optionally signed decimal integer, whose format is the same as expected for the subject sequence of `strtol(3C)` with the value 10 for the *base* argument. In the absence of a size modifier, the corresponding argument must be a pointer to `int`.
 - i Matches an optionally signed integer, whose format is the same as expected for the subject sequence of `strtol()` with 0 for the *base* argument. In the absence of a size modifier, the corresponding argument must be a pointer to `int`.
 - o Matches an optionally signed octal integer, whose format is the same as expected for the subject sequence of `strtoul(3C)` with the value 8 for the *base* argument. In the absence of a size modifier, the corresponding argument must be a pointer to unsigned `int`.
 - u Matches an optionally signed decimal integer, whose format is the same as expected for the subject sequence of `strtoul()` with the value 10 for the *base* argument. In the absence of a size modifier, the corresponding argument must be a pointer to unsigned `int`.
 - x Matches an optionally signed hexadecimal integer, whose format is the same as expected for the subject sequence of `strtoul()` with the value 16 for the *base* argument. In the absence of a size modifier, the corresponding argument must be a pointer to unsigned `int`.
 - a,e,f,g* Matches an optionally signed floating-point number, infinity, or NaN, whose format is the same as expected for the subject sequence of `strtod(3C)`. In the absence of a size modifier, the corresponding argument must be a pointer to `float`. The *e*, *f*, and *g* specifiers match hexadecimal floating point values only in C99/SUSv3 (see `standards(5)`) mode, but the *a* specifier always matches hexadecimal floating point values.
- These conversion specifiers match any subject sequence accepted by `strtod(3C)`, including the `INF`, `INFINITY`, `NAN`, and `NAN(n-char-sequence)` forms. The result of the conversion is the same as that of calling `strtod()` (or `strtodf()` or `strtold()`) with the matching sequence, including the raising of floating point exceptions and the setting of `errno` to `ERANGE`, if applicable.

s Matches a sequence of bytes that are not white-space characters. The corresponding argument must be a pointer to the initial byte of an array of `char`, `signed char`, or `unsigned char` large enough to accept the sequence and a terminating null character code, which will be added automatically.

If an `l` (`ell`) qualifier is present, the input is a sequence of characters that begins in the initial shift state. Each character is converted to a wide-character as if by a call to the `mbrtowc(3C)` function, with the conversion state described by an `mbstate_t` object initialized to zero before the first character is converted. The corresponding argument must be a pointer to an array of `wchar_t` large enough to accept the sequence and the terminating null wide-character, which will be added automatically.

[Matches a non-empty sequence of characters from a set of expected characters (the *scanset*). The normal skip over white-space characters is suppressed in this case. The corresponding argument must be a pointer to the initial byte of an array of `char`, `signed char`, or `unsigned char` large enough to accept the sequence and a terminating null byte, which will be added automatically.

If an `l` (`ell`) qualifier is present, the input is a sequence of characters that begins in the initial shift state. Each character in the sequence is converted to a wide-character as if by a call to the `mbrtowc()` function, with the conversion state described by an `mbstate_t` object initialized to zero before the first character is converted. The corresponding argument must be a pointer to an array of `wchar_t` large enough to accept the sequence and the terminating null wide-character, which will be added automatically.

The conversion specification includes all subsequent characters in the *format* string up to and including the matching right square bracket (`]`). The characters between the square brackets (the *scanlist*) comprise the *scanset*, unless the character after the left square bracket is a circumflex (`^`), in which case the *scanset* contains all characters that do not appear in the *scanlist* between the circumflex and the right square bracket. If the conversion specification begins with `[]` or `[^]`, the right square bracket is included in the *scanlist* and the next right square bracket is the matching right square bracket that ends the conversion specification; otherwise the first right square bracket is the one that ends the conversion specification. If a `-` is in the *scanlist* and is not the first character, nor the second where the first character is a `^`, nor the last character, it indicates a range of characters to be matched.

scanf(3C)

- | | |
|---|---|
| c | <p>Matches a sequence of characters of the number specified by the field width (1 if no field width is present in the conversion specification). The corresponding argument must be a pointer to the initial byte of an array of <code>char</code>, <code>signed char</code>, or <code>unsigned char</code> large enough to accept the sequence. No null byte is added. The normal skip over white-space characters is suppressed in this case.</p> <p>If an <code>l</code> (<code>ell</code>) qualifier is present, the input is a sequence of characters that begins in the initial shift state. Each character in the sequence is converted to a wide-character as if by a call to the <code>mbrtowc()</code> function, with the conversion state described by an <code>mbstate_t</code> object initialized to zero before the first character is converted. The corresponding argument must be a pointer to an array of <code>wchar_t</code> large enough to accept the resulting sequence of wide-characters. No null wide-character is added.</p> |
| p | <p>Matches the set of sequences that is the same as the set of sequences that is produced by the <code>%p</code> conversion of the corresponding <code>printf(3C)</code> functions. The corresponding argument must be a pointer to a pointer to <code>void</code>. If the input item is a value converted earlier during the same program execution, the pointer that results will compare equal to that value; otherwise the behavior of the <code>%p</code> conversion is undefined.</p> |
| n | <p>No input is consumed. The corresponding argument must be a pointer to the integer into which is to be written the number of bytes read from the input so far by this call to the <code>scanf()</code> functions. Execution of a <code>%n</code> conversion specification does not increment the assignment count returned at the completion of execution of the function.</p> |
| C | <p>Same as <code>lc</code>.</p> |
| S | <p>Same as <code>ls</code>.</p> |
| % | <p>Matches a single <code>%</code>; no conversion or assignment occurs. The complete conversion specification must be <code>%%</code>.</p> |

If a conversion specification is invalid, the behavior is undefined.

The conversion characters `A`, `E`, `F`, `G`, and `X` are also valid and behave the same as, respectively, `a`, `e`, `f`, `g`, and `x`.

If end-of-file is encountered during input, conversion is terminated. If end-of-file occurs before any bytes matching the current conversion specification (except for `%n`) have been read (other than leading white-space characters, where permitted), execution of the current conversion specification terminates with an input failure. Otherwise, unless execution of the current conversion specification is terminated with a matching failure, execution of the following conversion specification (if any) is terminated with an input failure.

Reaching the end of the string in `sscanf()` is equivalent to encountering end-of-file for `fscanf()`.

If conversion terminates on a conflicting input, the offending input is left unread in the input. Any trailing white space (including newline characters) is left unread unless matched by a conversion specification. The success of literal matches and suppressed assignments is only directly determinable via the `%n` conversion specification.

The `fscanf()` and `scanf()` functions may mark the `st_atime` field of the file associated with *stream* for update. The `st_atime` field will be marked for update by the first successful execution of `fgetc(3C)`, `fgets(3C)`, `fread(3C)`, `fscanf()`, `getc(3C)`, `getchar(3C)`, `gets(3C)`, or `scanf()` using *stream* that returns data not supplied by a prior call to `ungetc(3C)`.

RETURN VALUES

Upon successful completion, these functions return the number of successfully matched and assigned input items; this number can be 0 in the event of an early matching failure. If the input ends before the first matching failure or conversion, EOF is returned. If a read error occurs the error indicator for the stream is set, EOF is returned, and `errno` is set to indicate the error.

ERRORS

For the conditions under which the `scanf()` functions will fail and may fail, refer to `fgetc(3C)` or `fgetwc(3C)`.

In addition, `fscanf()` may fail if:

`EILSEQ` Input byte sequence does not form a valid character.
`EINVAL` There are insufficient arguments.

USAGE

If the application calling the `scanf()` functions has any objects of type `wint_t` or `wchar_t`, it must also include the header `<wchar.h>` to have these objects defined.

EXAMPLES

EXAMPLE 1 The call:

```
int i, n; float x; char name[50];
n = scanf("%d%f%s", &i, &x, name)
```

with the input line:

```
25 54.32E-1 Hamster
```

will assign to *n* the value 3, to *i* the value 25, to *x* the value 5.432, and *name* will contain the string Hamster.

EXAMPLE 2 The call:

```
int i; float x; char name[50];
(void) scanf("%2d%f*d %[0123456789]", &i, &x, name);
```

with input:

```
56789 0123 56a72
```

will assign 56 to *i*, 789.0 to *x*, skip 0123, and place the string 56\0 in *name*. The next call to `getchar(3C)` will return the character a.

scanf(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `fgetc(3C)`, `fgets(3C)`, `fgetwc(3C)`, `fread(3C)`, `isspace(3C)`, `printf(3C)`, `setlocale(3C)`, `strtod(3C)`, `strtol(3C)`, `strtoul(3C)`, `wcrtomb(3C)`, `ungetc(3C)`, `attributes(5)`, `standards(5)`

NAME	schedctl_init, schedctl_lookup, schedctl_exit, schedctl_start, schedctl_stop – preemption control
SYNOPSIS	<pre>cc [flag...] file... [library...] #include <schedctl.h> schedctl_t *schedctl_init(void); schedctl_t *schedctl_lookup(void); void schedctl_exit(void); void schedctl_start(schedctl_t *ptr); void schedctl_stop(schedctl_t *ptr);</pre>
DESCRIPTION	<p>These functions provide limited control over the scheduling of a thread (see threads(5)). They allow a running thread to give a hint to the kernel that preemptions of that thread should be avoided. The most likely use for these functions is to block preemption while holding a spinlock. Improper use of this facility, including attempts to block preemption for sustained periods of time, may result in reduced performance.</p> <p>The <code>schedctl_init()</code> function initializes preemption control for the calling thread and returns a pointer used to refer to the data. If <code>schedctl_init()</code> is called more than once by the same thread, the most recently returned pointer is the only valid one.</p> <p>The <code>schedctl_lookup()</code> function returns the currently allocated preemption control data associated with the calling thread that was previously returned by <code>schedctl_init()</code>. This can be useful in programs where it is difficult to maintain local state for each thread.</p> <p>The <code>schedctl_exit()</code> function removes the preemption control data associated with the calling thread.</p> <p>The <code>schedctl_start()</code> macro gives a hint to the kernel scheduler that preemption should be avoided on the current thread. The pointer passed to the macro must be the same as the pointer returned by the call to <code>schedctl_init()</code> by the current thread. The behavior of the program when other values are passed is undefined.</p> <p>The <code>schedctl_stop()</code> macro removes the hint that was set by <code>schedctl_start()</code>. As with <code>schedctl_start()</code>, the pointer passed to the macro must be the same as the pointer returned by the call to <code>schedctl_init()</code> by the current thread.</p> <p>The <code>schedctl_start()</code> and <code>schedctl_stop()</code> macros are intended to be used to bracket short critical sections, such as the time spent holding a spinlock. Other uses, including the failure to call <code>schedctl_stop()</code> soon after calling <code>schedctl_start()</code>, might result in poor performance.</p>
RETURN VALUES	The <code>schedctl_init()</code> function returns a pointer to a <code>schedctl_t</code> structure if the initialization was successful, or <code>NULL</code> otherwise. The <code>schedctl_lookup()</code> function returns a pointer to a <code>schedctl_t</code> structure if the data for that thread was found, or <code>NULL</code> otherwise.

schedctl_init(3C)

ERRORS No errors are returned.

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

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

SEE ALSO `prctl(1)`, `exec(2)`, `fork(2)`, `prctl(2)`, `attributes(5)`, `threads(5)`

NOTES Preemption control is intended for use by threads belonging to the time-sharing (TS), interactive (IA), fair-share (FSS), and fixed-priority (FX) scheduling classes. If used by threads in other scheduling classes, such as real-time (RT), no errors will be returned but `schedctl_start()` and `schedctl_stop()` will not have any effect.

The data used for preemption control are not copied in the child of a `fork(2)`. Thus, if a process containing threads using preemption control calls `fork` and the child does not immediately call `exec(2)`, each thread in the child must call `schedctl_init()` again prior to any future uses of `schedctl_start()` and `schedctl_stop()`. Failure to do so will result in undefined behavior.

NAME	seekdir – set position of directory stream						
SYNOPSIS	<pre>#include <sys/types.h> #include <dirent.h> void seekdir(DIR *dirp, long int loc);</pre>						
DESCRIPTION	<p>The <code>seekdir()</code> function sets the position of the next <code>readdir(3C)</code> operation on the directory stream specified by <i>dirp</i> to the position specified by <i>loc</i>. The value of <i>loc</i> should have been returned from an earlier call to <code>tellldir(3C)</code>. The new position reverts to the one associated with the directory stream when <code>tellldir()</code> was performed.</p> <p>If the value of <i>loc</i> was not obtained from an earlier call to <code>tellldir()</code> or if a call to <code>rewinddir(3C)</code> occurred between the call to <code>tellldir()</code> and the call to <code>seekdir()</code>, the results of subsequent calls to <code>readdir()</code> are unspecified.</p>						
RETURN VALUES	The <code>seekdir()</code> function returns no value.						
ERRORS	No errors are defined.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	<code>opendir(3C)</code> , <code>readdir(3C)</code> , <code>rewinddir(3C)</code> , <code>tellldir(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

select(3C)

NAME	select, pselect, FD_SET, FD_CLR, FD_ISSET, FD_ZERO – synchronous I/O multiplexing
SYNOPSIS	<pre>#include <sys/time.h> int select(int <i>nfds</i>, fd_set *restrict <i>readfds</i>, fd_set *restrict <i>writefds</i>, fd_set *restrict <i>errorfds</i>, struct timeval *restrict <i>timeout</i>); int pselect(int <i>nfds</i>, fd_set *restrict <i>readfds</i>, fd_set *restrict <i>writefds</i>, fd_set *restrict <i>errorfds</i>, const struct timespec *restrict <i>timeout</i>, const sigset_t *restrict <i>sigmask</i>); void FD_SET(int <i>fd</i>, fd_set *<i>fdset</i>); void FD_CLR(int <i>fd</i>, fd_set *<i>fdset</i>); int FD_ISSET(int <i>fd</i>, fd_set *<i>fdset</i>); void FD_ZERO(fd_set *<i>fdset</i>);</pre>
DESCRIPTION	<p>The <code>pselect()</code> function examines the file descriptor sets whose addresses are passed in the <code>readfds</code>, <code>writefds</code>, and <code>errorfds</code> parameters to see if some of their descriptors are ready for reading, are ready for writing, or have an exceptional condition pending, respectively.</p> <p>The <code>select()</code> function is equivalent to the <code>pselect()</code> function, except as follows:</p> <ul style="list-style-type: none">■ For the <code>select()</code> function, the timeout period is given in seconds and microseconds in an argument of type <code>struct timeval</code>, whereas for the <code>pselect()</code> function the timeout period is given in seconds and nanoseconds in an argument of type <code>struct timespec</code>.■ The <code>select()</code> function has no <code>sigmask</code> argument. It behaves as <code>pselect()</code> does when <code>sigmask</code> is a null pointer.■ Upon successful completion, the <code>select()</code> function might modify the object pointed to by the <code>timeout</code> argument. <p>The <code>select()</code> and <code>pselect()</code> functions support regular files, terminal and pseudo-terminal devices, STREAMS-based files, FIFOs, pipes, and sockets. The behavior of <code>select()</code> and <code>pselect()</code> on file descriptors that refer to other types of file is unspecified.</p> <p>The <code>nfds</code> argument specifies the range of file descriptors to be tested. The first <code>nfds</code> descriptors are checked in each set; that is, the descriptors from zero through <code>nfds-1</code> in the descriptor sets are examined.</p> <p>If the <code>readfds</code> argument is not a null pointer, it points to an object of type <code>fd_set</code> that on input specifies the file descriptors to be checked for being ready to read, and on output indicates which file descriptors are ready to read.</p> <p>If the <code>writefds</code> argument is not a null pointer, it points to an object of type <code>fd_set</code> that on input specifies the file descriptors to be checked for being ready to write, and on output indicates which file descriptors are ready to write.</p>

If the *errorfds* argument is not a null pointer, it points to an object of type `fd_set` that on input specifies the file descriptors to be checked for error conditions pending, and on output indicates which file descriptors have error conditions pending.

Upon successful completion, the objects pointed to by the *readfs*, *writes*, and *errorfds* arguments are modified to indicate which file descriptors are ready for reading, ready for writing, or have an error condition pending, respectively, and return the total number of ready descriptors in all the output sets. For each file descriptor less than *nfds*, the corresponding bit will be set on successful completion if it was set on input and the associated condition is true for that file descriptor.

If none of the selected descriptors are ready for the requested operation, the `select()` or `pselect()` function blocks until at least one of the requested operations becomes ready, until the timeout occurs, or until interrupted by a signal. The *timeout* parameter controls how long the `select()` or `pselect()` function takes before timing out. If the *timeout* parameter is not a null pointer, it specifies a maximum interval to wait for the selection to complete. If the specified time interval expires without any requested operation becoming ready, the function returns. If the *timeout* parameter is a null pointer, then the call to `select()` or `pselect()` blocks indefinitely until at least one descriptor meets the specified criteria. To effect a poll, the *timeout* parameter should not be a null pointer, and should point to a zero-valued `timespec` structure.

The use of a *timeout* does not affect any pending timers set up by `alarm(2)`, `ualarm(3C)`, or `setitimer(2)`.

If *sigmask* is not a null pointer, then the `pselect()` function replaces the signal mask of the process by the set of signals pointed to by *sigmask* before examining the descriptors, and restores the signal mask of the process before returning.

A descriptor is considered ready for reading when a call to an input function with `O_NONBLOCK` clear would not block, whether or not the function would transfer data successfully. (The function might return data, an end-of-file indication, or an error other than one indicating that it is blocked, and in each of these cases the descriptor will be considered ready for reading.)

A descriptor is considered ready for writing when a call to an output function with `O_NONBLOCK` clear would not block, whether or not the function would transfer data successfully.

If a socket has a pending error, it is considered to have an exceptional condition pending. Otherwise, what constitutes an exceptional condition is file type-specific. For a file descriptor for use with a socket, it is protocol-specific except as noted below. For other file types, if the operation is meaningless for a particular file type, `select()` or `pselect()` indicates that the descriptor is ready for read or write operations and indicates that the descriptor has no exceptional condition pending.

If a descriptor refers to a socket, the implied input function is the `recvmsg(3XNET)` function with parameters requesting normal and ancillary data, such that the presence of either type causes the socket to be marked as readable. The presence of out-of-band

select(3C)

data is checked if the socket option `SO_OOBINLINE` has been enabled, as out-of-band data is enqueued with normal data. If the socket is currently listening, then it is marked as readable if an incoming connection request has been received, and a call to the `accept` function completes without blocking.

If a descriptor refers to a socket, the implied output function is the `sendmsg(3XNET)` function supplying an amount of normal data equal to the current value of the `SO_SNDLOWAT` option for the socket. If a non-blocking call to the `connect` function has been made for a socket, and the connection attempt has either succeeded or failed leaving a pending error, the socket is marked as writable.

A socket is considered to have an exceptional condition pending if a receive operation with `O_NONBLOCK` clear for the open file description and with the `MSG_OOB` flag set would return out-of-band data without blocking. (It is protocol-specific whether the `MSG_OOB` flag would be used to read out-of-band data.) A socket will also be considered to have an exceptional condition pending if an out-of-band data mark is present in the receive queue.

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

Selecting true for reading on a socket descriptor upon which a `listen(3XNET)` call has been performed indicates that a subsequent `accept(3XNET)` call on that descriptor will not block.

If the *timeout* argument is not a null pointer, it points to an object of type `struct timeval` that specifies a maximum interval to wait for the selection to complete. If the *timeout* argument points to an object of type `struct timeval` whose members are 0, `select()` does not block. If the *timeout* argument is a null pointer, `select()` blocks until an event causes one of the masks to be returned with a valid (non-zero) value. If the time limit expires before any event occurs that would cause one of the masks to be set to a non-zero value, `select()` completes successfully and returns 0.

If the *readfs*, *writes*, and *errorfds* arguments are all null pointers and the *timeout* argument is not a null pointer, `select()` or `pselect()` blocks for the time specified, or until interrupted by a signal. If the *readfs*, *writes*, and *errorfds* arguments are all null pointers and the *timeout* argument is a null pointer, `select()` blocks until interrupted by a signal.

File descriptors associated with regular files always select true for ready to read, ready to write, and error conditions.

On failure, the objects pointed to by the *readfs*, *writes*, and *errorfds* arguments are not modified. If the timeout interval expires without the specified condition being true for any of the specified file descriptors, the objects pointed to by the *readfs*, *writes*, and *errorfds* arguments have all bits set to 0.

File descriptor masks of type `fd_set` can be initialized and tested with the macros `FD_CLR()`, `FD_ISSET()`, `FD_SET()`, and `FD_ZERO()`.

<code>FD_CLR(<i>fd</i>, &<i>fdset</i>)</code>	Clears the bit for the file descriptor <i>fd</i> in the file descriptor set <i>fdset</i> .
<code>FD_ISSET(<i>fd</i>, &<i>fdset</i>)</code>	Returns a non-zero value if the bit for the file descriptor <i>fd</i> is set in the file descriptor set pointed to by <i>fdset</i> , and 0 otherwise.
<code>FD_SET(<i>fd</i>, &<i>fdset</i>)</code>	Sets the bit for the file descriptor <i>fd</i> in the file descriptor set <i>fdset</i> .
<code>FD_ZERO(&<i>fdset</i>)</code>	Initializes the file descriptor set <i>fdset</i> to have zero bits for all file descriptors.

The behavior of these macros is undefined if the *fd* argument is less than 0 or greater than or equal to `FD_SETSIZE`, or if *fd* is not a valid file descriptor, or if any of the arguments are expressions with side effects.

RETURN VALUES

On successful completion, `select()` and `pselect()` return the total number of bits set in the bit masks. Otherwise, `-1` is returned and `errno` is set to indicate the error.

The `FD_CLR()`, `FD_SET()`, and `FD_ZERO()` macros return no value. The `FD_ISSET()` macro returns a non-zero value if the bit for the file descriptor *fd* is set in the file descriptor set pointed to by *fdset*, and 0 otherwise.

ERRORS

The `select()` and `pselect()` functions will fail if:

<code>EBADF</code>	One or more of the file descriptor sets specified a file descriptor that is not a valid open file descriptor.
<code>EINTR</code>	The function was interrupted before any of the selected events occurred and before the timeout interval expired. If <code>SA_RESTART</code> has been set for the interrupting signal, it is implementation-dependent whether <code>select()</code> restarts or returns with <code>EINTR</code> .
<code>EINVAL</code>	An invalid timeout interval was specified.
<code>EINVAL</code>	The <i>nfds</i> argument is less than 0 or greater than <code>FD_SETSIZE</code> .
<code>EINVAL</code>	One of the specified file descriptors refers to a <code>STREAM</code> or multiplexer that is linked (directly or indirectly) downstream from a multiplexer.
<code>EINVAL</code>	A component of the pointed-to time limit is outside the acceptable range: <code>t_sec</code> must be between 0 and 10^8 , inclusive. <code>t_usec</code> must be greater than or equal to 0, and less than 10^6 .

USAGE

The `poll(2)` function is preferred over this function. It must be used when the number of file descriptors exceeds `FD_SETSIZE`.

The use of a timeout does not affect any pending timers set up by `alarm(2)`, `ualarm(3C)` or `setitimer(2)`.

select(3C)

On successful completion, the object pointed to by the *timeout* argument may be modified.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `alarm(2)`, `fcntl(2)`, `poll(2)`, `read(2)`, `setitimer(2)`, `write(2)`, `accept(3SOCKET)`, `listen(3SOCKET)`, `ualarm(3C)`, `attributes(5)`, `standards(5)`

NOTES The default value for `FD_SETSIZE` (currently 1024) is larger than the default limit on the number of open files. To accommodate 32-bit applications that wish to use a larger number of open files with `select()`, it is possible to increase this size at compile time by providing a larger definition of `FD_SETSIZE` before the inclusion of any system-supplied header. The maximum supported size for `FD_SETSIZE` is 65536. The default value is already 65536 for 64-bit applications.

NAME	semaphore, sema_init, sema_destroy, sema_wait, sema_trywait, sema_post – semaphores				
SYNOPSIS	<pre>cc [<i>flag...</i>] <i>file...</i> -lthread -lc [<i>library...</i>] #include <synch.h> int sema_init(sema_t *sp, unsigned int <i>count</i>, int <i>type</i>, void * <i>arg</i>); int sema_destroy(sema_t *sp); int sema_wait(sema_t *sp); int sema_trywait(sema_t *sp); int sema_post(sema_t *sp);</pre>				
DESCRIPTION	<p>A semaphore is a non-negative integer count and is generally used to coordinate access to resources. The initial semaphore count is set to the number of free resources, then threads slowly increment and decrement the count as resources are added and removed. If the semaphore count drops to 0, which means no available resources, threads attempting to decrement the semaphore will block until the count is greater than 0.</p> <p>Semaphores can synchronize threads in this process and other processes if they are allocated in writable memory and shared among the cooperating processes (see <code>mmap(2)</code>), and have been initialized for this purpose.</p> <p>Semaphores must be initialized before use; semaphores pointed to by <i>sp</i> to <i>count</i> are initialized by <code>sema_init()</code>. The <i>type</i> argument can assign several different types of behavior to a semaphore. No current type uses <i>arg</i>, although it may be used in the future.</p> <p>The <i>type</i> argument may be one of the following:</p> <table border="0" style="width: 100%;"> <tr> <td style="vertical-align: top; padding-right: 20px;">USYNC_PROCESS</td> <td>The semaphore can synchronize threads in this process and other processes. Initializing the semaphore should be done by only one process. A semaphore initialized with this type must be allocated in memory shared between processes, either in Sys V shared memory (see <code>shmop(2)</code>), or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and not allocate it in such shared memory. <i>arg</i> is ignored.</td> </tr> <tr> <td style="vertical-align: top; padding-right: 20px;">USYNC_THREAD</td> <td>The semaphore can synchronize threads only in this process. The <i>arg</i> argument is ignored. USYNC_THREAD does not support multiple mappings to the same logical synch object. If you need to <code>mmap()</code> a synch object to different locations within the same address space, then the synch object should be initialized as a shared object USYNC_PROCESS for Solaris threads and PTHREAD_PROCESS_PRIVATE for POSIX threads.</td> </tr> </table>	USYNC_PROCESS	The semaphore can synchronize threads in this process and other processes. Initializing the semaphore should be done by only one process. A semaphore initialized with this type must be allocated in memory shared between processes, either in Sys V shared memory (see <code>shmop(2)</code>), or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and not allocate it in such shared memory. <i>arg</i> is ignored.	USYNC_THREAD	The semaphore can synchronize threads only in this process. The <i>arg</i> argument is ignored. USYNC_THREAD does not support multiple mappings to the same logical synch object. If you need to <code>mmap()</code> a synch object to different locations within the same address space, then the synch object should be initialized as a shared object USYNC_PROCESS for Solaris threads and PTHREAD_PROCESS_PRIVATE for POSIX threads.
USYNC_PROCESS	The semaphore can synchronize threads in this process and other processes. Initializing the semaphore should be done by only one process. A semaphore initialized with this type must be allocated in memory shared between processes, either in Sys V shared memory (see <code>shmop(2)</code>), or in memory mapped to a file (see <code>mmap(2)</code>). It is illegal to initialize the object this way and not allocate it in such shared memory. <i>arg</i> is ignored.				
USYNC_THREAD	The semaphore can synchronize threads only in this process. The <i>arg</i> argument is ignored. USYNC_THREAD does not support multiple mappings to the same logical synch object. If you need to <code>mmap()</code> a synch object to different locations within the same address space, then the synch object should be initialized as a shared object USYNC_PROCESS for Solaris threads and PTHREAD_PROCESS_PRIVATE for POSIX threads.				

semaphore(3C)

A semaphore must not be simultaneously initialized by multiple threads, nor re-initialized while in use by other threads.

Default semaphore initialization (intra-process):

```
sema_t sp;
int count = 1;
sema_init(&sp, count, NULL, NULL);
```

or

```
sema_init(&sp, count, USYNC_THREAD, NULL);
```

Customized semaphore initialization (inter-process):

```
sema_t sp;
int count = 1;
sema_init(&sp, count, USYNC_PROCESS, NULL);
```

The `sema_destroy()` function destroys any state related to the semaphore pointed to by `sp`. The semaphore storage space is not released.

The `sema_wait()` function blocks the calling thread until the semaphore count pointed to by `sp` is greater than 0, and then it atomically decrements the count.

The `sema_trywait()` function atomically decrements the semaphore count pointed to by `sp`, if the count is greater than 0; otherwise, it returns an error.

The `sema_post()` function atomically increments the semaphore count pointed to by `sp`. If there are any threads blocked on the semaphore, one will be unblocked.

The semaphore functionality described on this man page is for the Solaris threads implementation. For the POSIX-conforming semaphore interface documentation, see `sem_close(3RT)`, `sem_destroy(3RT)`, `sem_getvalue(3RT)`, `sem_init(3RT)`, `sem_open(3RT)`, `sem_post(3RT)`, `sem_unlink(3RT)`, and `sem_wait(3RT)`.

RETURN VALUES

Upon successful completion, 0 is returned; otherwise, a non-zero value indicates an error.

ERRORS

These functions will fail if:

`EINVAL` The `sp` argument does not refer to a valid semaphore.

`EFAULT` Either the `sp` or `arg` argument points to an illegal address.

The `sema_wait()` function will fail if:

`EINTR` The wait was interrupted by a signal or `fork()`.

The `sema_trywait()` function will fail if:

`EBUSY` The semaphore pointed to by `sp` has a 0 count.

The `sema_post()` function will fail if:

EOVERFLOW The semaphore value pointed to by *sp* exceeds SEM_VALUE_MAX.

EXAMPLES

EXAMPLE 1 The customer waiting-line in a bank is analogous to the synchronization scheme of a semaphore using `sema_wait()` and `sema_trywait()`:

```
/* cc [ flag . . . ] file . . . -lthread [ library . . . ] */
#include <errno.h>
#define TELLERS 10
sema_t    tellers;    /* semaphore */
int banking_hours(), deposit_withdrawal;
void*customer(), do_business(), skip_banking_today();
. . .

sema_init(&tellers, TELLERS, USYNC_THREAD, NULL);
    /* 10 tellers available */
while(banking_hours())
    pthread_create(NULL, NULL, customer, deposit_withdrawal);
. . .

void *
customer(int deposit_withdrawal)
{
    int this_customer, in_a_hurry = 50;
    this_customer = rand() % 100;

    if (this_customer == in_a_hurry) {
        if (sema_trywait(&tellers) != 0)
            if (errno == EBUSY){ /* no teller available */
                skip_banking_today(this_customer);
                return;
            } /* else go immediately to available teller and
                decrement tellers */
    }
    else
        sema_wait(&tellers); /* wait for next teller, then proceed,
                            and decrement tellers */

    do_business(deposit_withdrawal);
    sema_post(&tellers); /* increment tellers;
                        this_customer's teller
                        is now available */
}
```

ATTRIBUTES

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

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

SEE ALSO

`mmap(2)`, `shmop(2)`, `sem_close(3RT)`, `sem_destroy(3RT)`, `sem_getvalue(3RT)`, `sem_init(3RT)`, `sem_open(3RT)`, `sem_post(3RT)`, `sem_unlink(3RT)`, `sem_wait(3RT)`, `attributes(5)`, `standards(5)`

NOTES

These functions are also available by way of:

semaphore(3C)

```
#include <thread.h>
```

By default, there is no defined order of unblocking for multiple threads waiting for a semaphore.

NAME	setbuf, setvbuf – assign buffering to a stream						
SYNOPSIS	<pre>#include <stdio.h> void setbuf(FILE *stream, char *buf); int setvbuf(FILE *stream, char *buf, int type, size_t size);</pre>						
DESCRIPTION	<p>The <code>setbuf()</code> function may be used after the stream pointed to by <code>stream</code> (see intro(3)) is opened but before it is read or written. It causes the array pointed to by <code>buf</code> to be used instead of an automatically allocated buffer. If <code>buf</code> is the null pointer, input/output will be completely unbuffered. The constant <code>BUFSIZ</code>, defined in the <code><stdio.h></code> header, indicates the size of the array pointed to by <code>buf</code>.</p> <p>The <code>setvbuf()</code> function may be used after a stream is opened but before it is read or written. The <code>type</code> argument determines how <code>stream</code> will be buffered. Legal values for <code>type</code> (defined in <code><stdio.h></code>) are:</p> <table border="0"> <tr> <td style="padding-right: 20px;"><code>_IOFBF</code></td> <td>Input/output to be fully buffered.</td> </tr> <tr> <td><code>_IOLBF</code></td> <td>Output to be line buffered; the buffer will be flushed when a <code>NEWLINE</code> is written, the buffer is full, or input is requested.</td> </tr> <tr> <td><code>_IONBF</code></td> <td>Input/output to be completely unbuffered.</td> </tr> </table> <p>If <code>buf</code> is not the null pointer, the array it points to will be used for buffering, instead of an automatically allocated buffer. The <code>size</code> argument specifies the size of the buffer to be used. If input/output is unbuffered, <code>buf</code> and <code>size</code> are ignored.</p> <p>For a further discussion of buffering, see stdio(3C).</p>	<code>_IOFBF</code>	Input/output to be fully buffered.	<code>_IOLBF</code>	Output to be line buffered; the buffer will be flushed when a <code>NEWLINE</code> is written, the buffer is full, or input is requested.	<code>_IONBF</code>	Input/output to be completely unbuffered.
<code>_IOFBF</code>	Input/output to be fully buffered.						
<code>_IOLBF</code>	Output to be line buffered; the buffer will be flushed when a <code>NEWLINE</code> is written, the buffer is full, or input is requested.						
<code>_IONBF</code>	Input/output to be completely unbuffered.						
RETURN VALUES	If an illegal value for <code>type</code> is provided, <code>setvbuf()</code> returns a non-zero value. Otherwise, it returns 0.						
USAGE	<p>A common source of error is allocating buffer space as an “automatic” variable in a code block, and then failing to close the stream in the same block.</p> <p>When using <code>setbuf()</code>, <code>buf</code> should always be sized using <code>BUFSIZ</code>. If the array pointed to by <code>buf</code> is larger than <code>BUFSIZ</code>, a portion of <code>buf</code> will not be used. If <code>buf</code> is smaller than <code>BUFSIZ</code>, other memory may be unexpectedly overwritten.</p> <p>Parts of <code>buf</code> will be used for internal bookkeeping of the stream and, therefore, <code>buf</code> will contain less than <code>size</code> bytes when full. It is recommended that stdio(3C) be used to handle buffer allocation when using <code>setvbuf()</code>.</p>						
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:						

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard

setbuf(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [fopen\(3C\)](#), [getc\(3C\)](#), [malloc\(3C\)](#), [putc\(3C\)](#), [stdio\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	setbuffer, setlinebuf – assign buffering to a stream
SYNOPSIS	<pre>#include <stdio.h> void setbuffer(FILE *iop, char *abuf, size_t asize); int setlinebuf(FILE *iop);</pre>
DESCRIPTION	<p>The <code>setbuffer()</code> and <code>setlinebuf()</code> functions assign buffering to a stream. The three types of buffering available are unbuffered, block buffered, and line buffered. When an output stream is unbuffered, information appears on the destination file or terminal as soon as written; when it is block buffered, many characters are saved and written as a block; when it is line buffered, characters are saved until either a NEWLINE is encountered or input is read from <code>stdin</code>. The <code>fflush(3C)</code> function may be used to force the block out early. Normally all files are block buffered. A buffer is obtained from <code>malloc(3C)</code> upon the first <code>getc(3C)</code> or <code>putc(3C)</code> performed on the file. If the standard stream <code>stdout</code> refers to a terminal, it is line buffered. The standard stream <code>stderr</code> is unbuffered by default.</p> <p>The <code>setbuffer()</code> function can be used after a stream <code>iop</code> has been opened but before it is read or written. It uses the character array <code>abuf</code> whose size is determined by the <code>asize</code> argument instead of an automatically allocated buffer. If <code>abuf</code> is the null pointer, input/output will be completely unbuffered. A manifest constant <code>BUFSIZ</code>, defined in the <code><stdio.h></code> header, tells how large an array is needed:</p> <pre>char buf[BUFSIZ];</pre> <p>The <code>setlinebuf()</code> function is used to change the buffering on a stream from block buffered or unbuffered to line buffered. Unlike <code>setbuffer()</code>, it can be used at any time that the stream <code>iop</code> is active.</p> <p>A stream can be changed from unbuffered or line buffered to block buffered by using <code>freopen(3C)</code>. A stream can be changed from block buffered or line buffered to unbuffered by using <code>freopen(3C)</code> followed by <code>setbuf(3C)</code> with a buffer argument of <code>NULL</code>.</p>
RETURN VALUES	The <code>setlinebuf()</code> function returns no useful value.
SEE ALSO	<code>malloc(3C)</code> , <code>fclose(3C)</code> , <code>fopen(3C)</code> , <code>fread(3C)</code> , <code>getc(3C)</code> , <code>printf(3C)</code> , <code>putc(3C)</code> , <code>puts(3C)</code> , <code>setbuf(3C)</code> , <code>setvbuf(3C)</code>
NOTES	A common source of error is allocating buffer space as an “automatic” variable in a code block, and then failing to close the stream in the same block.

setcat(3C)

NAME setcat – define default catalog

SYNOPSIS

```
#include <pfmt.h>

char *setcat(const char *catalog);
```

DESCRIPTION The `setcat()` function defines the default message catalog to be used by subsequent calls to `gettext(3C)`, `lfmt(3C)`, or `pfmt(3C)` that do not explicitly specify a message catalog.

The *catalog* argument must be limited to 14 characters. These characters must be selected from a set of all characters values, excluding `\0` (null) and the ASCII codes for `/` (slash) and `:` (colon).

The `setcat()` function assumes that the catalog exists. No checking is done on the argument.

A null pointer passed as an argument will result in the return of a pointer to the current default message catalog name. A pointer to an empty string passed as an argument will cancel the default catalog.

If no default catalog is specified, or if *catalog* is an invalid catalog name, subsequent calls to `gettext(3C)`, `lfmt(3C)`, or `pfmt(3C)` that do not explicitly specify a catalog name will use `Message not found!!\n` as default string.

RETURN VALUES Upon successful completion, `setcat()` returns a pointer to the catalog name. Otherwise, it returns a null pointer.

EXAMPLES **EXAMPLE 1** Example of `setcat()` function.

```
setcat("test");
gettext(":10", "hello world\n")
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `gettext(3C)`, `lfmt(3C)`, `pfmt(3C)`, `setlocale(3C)`, `attributes(5)`, `environ(5)`

NAME	setenv – add or change environment variable						
SYNOPSIS	<pre>#include <stdlib.h> int setenv(const char *<i>envname</i>, const char *<i>envval</i>, int <i>overwrite</i>);</pre>						
DESCRIPTION	<p>The <code>setenv()</code> function updates or adds a variable in the environment of the calling process. The <i>envname</i> argument points to a string containing the name of an environment variable to be added or altered. The environment variable is set to the value to which <i>envval</i> points. The function fails if <i>envname</i> points to a string which contains an '=' character. If the environment variable named by <i>envname</i> already exists and the value of <i>overwrite</i> is non-zero, the function returns successfully and the environment is updated. If the environment variable named by <i>envname</i> already exists and the value of <i>overwrite</i> is zero, the function returns successfully and the environment remains unchanged.</p> <p>If the application modifies <i>environ</i> or the pointers to which it points, the behavior of <code>setenv()</code> is undefined. The <code>setenv()</code> function updates the list of pointers to which <i>environ</i> points.</p> <p>The strings described by <i>envname</i> and <i>envval</i> are copied by this function.</p>						
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, <code>errno</code> set to indicate the error, and the environment is left unchanged.						
ERRORS	<p>The <code>setenv()</code> function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EINVAL</td> <td>The <i>envname</i> argument is a null pointer, points to an empty string, or points to a string containing an '=' character.</td> </tr> <tr> <td>ENOMEM</td> <td>Insufficient memory was available to add a variable or its value to the environment.</td> </tr> </table>	EINVAL	The <i>envname</i> argument is a null pointer, points to an empty string, or points to a string containing an '=' character.	ENOMEM	Insufficient memory was available to add a variable or its value to the environment.		
EINVAL	The <i>envname</i> argument is a null pointer, points to an empty string, or points to a string containing an '=' character.						
ENOMEM	Insufficient memory was available to add a variable or its value to the environment.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>getenv(3C)</code> , <code>unsetenv(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

setjmp(3C)

NAME	setjmp, sigsetjmp, longjmp, siglongjmp – non-local goto
SYNOPSIS	<pre>#include <setjmp.h> int setjmp(jmp_buf <i>env</i>); int sigsetjmp(sigjmp_buf <i>env</i>, int <i>savemask</i>); void longjmp(jmp_buf <i>env</i>, int <i>val</i>); void siglongjmp(sigjmp_buf <i>env</i>, int <i>val</i>);</pre>
DESCRIPTION	<p>These functions are useful for dealing with errors and interrupts encountered in a low-level subroutine of a program.</p> <p>The <code>setjmp()</code> function saves its stack environment in <i>env</i> for later use by <code>longjmp()</code>.</p> <p>The <code>sigsetjmp()</code> function saves the calling process's registers and stack environment (see <code>sigaltstack(2)</code>) in <i>env</i> for later use by <code>siglongjmp()</code>. If <i>savemask</i> is non-zero, the calling process's signal mask (see <code>sigprocmask(2)</code>) and scheduling parameters (see <code>prctl(2)</code>) are also saved.</p> <p>The <code>longjmp()</code> function restores the environment saved by the last call of <code>setjmp()</code> with the corresponding <i>env</i> argument. After <code>longjmp()</code> completes, program execution continues as if the corresponding call to <code>setjmp()</code> had just returned the value <i>val</i>. The caller of <code>setjmp()</code> must not have returned in the interim. The <code>longjmp()</code> function cannot cause <code>setjmp()</code> to return the value 0. If <code>longjmp()</code> is invoked with a second argument of 0, <code>setjmp()</code> will return 1. At the time of the second return from <code>setjmp()</code>, all external and static variables have values as of the time <code>longjmp()</code> is called (see EXAMPLES).</p> <p>The <code>siglongjmp()</code> function restores the environment saved by the last call of <code>sigsetjmp()</code> with the corresponding <i>env</i> argument. After <code>siglongjmp()</code> completes, program execution continues as if the corresponding call to <code>sigsetjmp()</code> had just returned the value <i>val</i>. The <code>siglongjmp()</code> function cannot cause <code>sigsetjmp()</code> to return the value 0. If <code>siglongjmp()</code> is invoked with a second argument of 0, <code>sigsetjmp()</code> will return 1. At the time of the second return from <code>sigsetjmp()</code>, all external and static variables have values as of the time <code>siglongjmp()</code> was called.</p> <p>If a signal-catching function interrupts <code>sleep(3C)</code> and calls <code>siglongjmp()</code> to restore an environment saved prior to the <code>sleep()</code> call, the action associated with <code>SIGALRM</code> and time it is scheduled to be generated are unspecified. It is also unspecified whether the <code>SIGALRM</code> signal is blocked, unless the process's signal mask is restored as part of the environment.</p> <p>The <code>siglongjmp()</code> function restores the saved signal mask if and only if the <i>env</i> argument was initialized by a call to the <code>sigsetjmp()</code> function with a non-zero <i>savemask</i> argument.</p> <p>The values of register and automatic variables are undefined. Register or automatic variables whose value must be relied upon must be declared as volatile.</p>

RETURN VALUES If the return is from a direct invocation, `setjmp()` and `sigsetjmp()` return 0. If the return is from a call to `longjmp()`, `setjmp()` returns a non-zero value. If the return is from a call to `siglongjmp()`, `sigsetjmp()` returns a non-zero value.

After `longjmp()` is completed, program execution continues as if the corresponding invocation of `setjmp()` had just returned the value specified by *val*. The `longjmp()` function cannot cause `setjmp()` to return 0; if *val* is 0, `setjmp()` returns 1.

After `siglongjmp()` is completed, program execution continues as if the corresponding invocation of `sigsetjmp()` had just returned the value specified by *val*. The `siglongjmp()` function cannot cause `sigsetjmp()` to return 0; if *val* is 0, `sigsetjmp()` returns 1.

EXAMPLES **EXAMPLE 1** Example of `setjmp()` and `longjmp()` functions.

The following example uses both `setjmp()` and `longjmp()` to return the flow of control to the appropriate instruction block:

```
#include <stdio.h>
#include <setjmp.h>
#include <signal.h>
#include <unistd.h>
jmp_buf env; static void signal_handler();

main() {
    int returned_from_longjump, processing = 1;
    unsigned int time_interval = 4;
    if ((returned_from_longjump = setjmp(env)) != 0)
        switch (returned_from_longjump) {
            case SIGINT:
                printf("longjumped from interrupt %d\n", SIGINT);
                break;
            case SIGALRM:
                printf("longjumped from alarm %d\n", SIGALRM);
                break;
        }
    (void) signal(SIGINT, signal_handler);
    (void) signal(SIGALRM, signal_handler);
    alarm(time_interval);
    while (processing) {
        printf(" waiting for you to INTERRUPT (cntrl-C) ...\n");
        sleep(1);
    }
    /* end while forever loop */
}

static void signal_handler(sig)
int sig; {
    switch (sig) {
        case SIGINT:
            ... /* process for interrupt */
            longjmp(env, sig);
            /* break never reached */
        case SIGALRM:
            ... /* process for alarm */
            longjmp(env, sig);
            /* break never reached */
        default:
            exit(sig);
    }
}
```

setjmp(3C)

EXAMPLE 1 Example of `setjmp()` and `longjmp()` functions. (Continued)

When this example is compiled and executed, and the user sends an interrupt signal, the output will be:

```
longjumped from interrupt
```

Additionally, every 4 seconds the alarm will expire, signalling this process, and the output will be:

```
longjumped from alarm
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Unsafe

SEE ALSO `getcontext(2)`, `pricntl(2)`, `sigaction(2)`, `sigaltstack(2)`, `sigprocmask(2)`, `signal(3C)`, `attributes(5)`, `standards(5)`

WARNINGS If `longjmp()` or `siglongjmp()` are called even though `env` was never primed by a call to `setjmp()` or `sigsetjmp()`, or when the last such call was in a function that has since returned, the results are undefined.

NAME	setjmp, longjmp, _setjmp, _longjmp – non-local goto
SYNOPSIS	<pre> /usr/ucb/cc [flag ...] file ... #include <setjmp.h> int setjmp(env); jmp_buf env; void longjmp(env, val); jmp_buf env; int val; int _setjmp(env); jmp_buf env; void _longjmp(env, val); jmp_buf env; int val; </pre>
DESCRIPTION	<p>The <code>setjmp()</code> and <code>longjmp()</code> functions are useful for dealing with errors and interrupts encountered in a low-level subroutine of a program.</p> <p>The <code>setjmp()</code> function saves its stack environment in <code>env</code> for later use by <code>longjmp()</code>. A normal call to <code>setjmp()</code> returns zero. <code>setjmp()</code> also saves the register environment. If a <code>longjmp()</code> call will be made, the routine which called <code>setjmp()</code> should not return until after the <code>longjmp()</code> has returned control (see below).</p> <p>The <code>longjmp()</code> function restores the environment saved by the last call of <code>setjmp()</code>, and then returns in such a way that execution continues as if the call of <code>setjmp()</code> had just returned the value <code>val</code> to the function that invoked <code>setjmp()</code>; however, if <code>val</code> were zero, execution would continue as if the call of <code>setjmp()</code> had returned one. This ensures that a “return” from <code>setjmp()</code> caused by a call to <code>longjmp()</code> can be distinguished from a regular return from <code>setjmp()</code>. The calling function must not itself have returned in the interim, otherwise <code>longjmp()</code> will be returning control to a possibly non-existent environment. All memory-bound data have values as of the time <code>longjmp()</code> was called. The CPU and floating-point data registers are restored to the values they had at the time that <code>setjmp()</code> was called. But, because the <code>register</code> storage class is only a hint to the C compiler, variables declared as <code>register</code> variables may not necessarily be assigned to machine registers, so their values are unpredictable after a <code>longjmp()</code>. This is especially a problem for programmers trying to write machine-independent C routines.</p> <p>The <code>setjmp()</code> and <code>longjmp()</code> functions save and restore the signal mask while <code>_setjmp()</code> and <code>_longjmp()</code> manipulate only the C stack and registers.</p> <p>None of these functions save or restore any floating-point status or control registers.</p>

setjmp(3UCB)

EXAMPLES **EXAMPLE 1** Examples of `setjmp()` and `longjmp()`.

The following example uses both `setjmp()` and `longjmp()` to return the flow of control to the appropriate instruction block:

```
#include <stdio.h>
#include <setjmp.h>
#include <signal.h>
#include <unistd.h>
jmp_buf env; static void signal_handler();
main() {
    int returned_from_longjump, processing = 1;
    unsigned int time_interval = 4;
    if ((returned_from_longjump = setjmp(env)) != 0)
        switch (returned_from_longjump) {
            case SIGINT:
                printf("longjumped from interrupt %d\n",SIGINT);
                break;
            case SIGALRM:
                printf("longjumped from alarm %d\n",SIGALRM);
                break;
        }
    (void) signal(SIGINT, signal_handler);
    (void) signal(SIGALRM, signal_handler);
    alarm(time_interval);
    while (processing) {
        printf(" waiting for you to INTERRUPT (cntrl-C) ...\n");
        sleep(1);
    } /* end while forever loop */
}

static void signal_handler(sig)
int sig; {
    switch (sig) {
        case SIGINT:
            ... /* process for interrupt */
            longjmp(env, sig);
            /* break never reached */
        case SIGALRM:
            ... /* process for alarm */
            longjmp(env, sig);
            /* break never reached */
        default:
            exit(sig);
    }
}
```

When this example is compiled and executed, and the user sends an interrupt signal, the output will be:

```
longjumped from interrupt
```

Additionally, every 4 seconds the alarm will expire, signalling this process, and the output will be:

```
longjumped from alarm
```

SEE ALSO `cc(1B)`, `sigvec(3UCB)`, `setjmp(3C)`, `signal(3C)`

NOTES Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.

BUGS The `setjmp()` function does not save the current notion of whether the process is executing on the signal stack. The result is that a `longjmp()` to some place on the signal stack leaves the signal stack state incorrect.

On some systems `setjmp()` also saves the register environment. Therefore, all data that are bound to registers are restored to the values they had at the time that `setjmp()` was called. All memory-bound data have values as of the time `longjmp()` was called. However, because the `register` storage class is only a hint to the C compiler, variables declared as `register` variables may not necessarily be assigned to machine registers, so their values are unpredictable after a `longjmp()`. When using compiler options that specify automatic register allocation (see `cc(1B)`), the compiler will not attempt to assign variables to registers in routines that call `setjmp()`.

The `longjmp()` function never causes `setjmp()` to return 0, so programmers should not depend on `longjmp()` being able to cause `setjmp()` to return 0.

setkey(3C)

NAME	setkey – set encoding key						
SYNOPSIS	<pre>#include <stdlib.h> void setkey(const char *key);</pre>						
DESCRIPTION	The <code>setkey()</code> function provides (rather primitive) access to the hashing algorithm employed by the <code>crypt(3C)</code> function. The argument of <code>setkey()</code> is an array of length 64 bytes containing only the bytes with numerical value of 0 and 1. If this string is divided into groups of 8, the low-order bit in each group is ignored; this gives a 56-bit key which is used by the algorithm. This is the key that will be used with the algorithm to encode a string <i>block</i> passed to <code>encrypt(3C)</code> .						
RETURN VALUES	No values are returned.						
ERRORS	The <code>setkey()</code> function will fail if: ENOSYS The functionality is not supported on this implementation.						
USAGE	In some environments, decoding may not be implemented. This is related to U.S. Government restrictions on encryption and decryption routines: the DES decryption algorithm cannot be exported outside the U.S.A. Historical practice has been to ship a different version of the encryption library without the decryption feature in the routines supplied. Thus the exported version of <code>encrypt()</code> does encoding but not decoding. Because <code>setkey()</code> does not return a value, applications wishing to check for errors should set <code>errno</code> to 0, call <code>setkey()</code> , then test <code>errno</code> and, if it is non-zero, assume an error has occurred.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes: <table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	<code>crypt(3C)</code> , <code>encrypt(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	setlabel – define the label for <code>pfmt()</code> and <code>lfmt()</code>				
SYNOPSIS	<pre>#include <pfmt..h> int setlabel(const char *label);</pre>				
DESCRIPTION	<p>The <code>setlabel()</code> function defines the label for messages produced in standard format by subsequent calls to <code>lfmt(3C)</code> and <code>pfmt(3C)</code>.</p> <p>The <i>label</i> argument is a character string no more than 25 characters in length.</p> <p>No label is defined before <code>setlabel()</code> is called. The label should be set once at the beginning of a utility and remain constant. A null pointer or an empty string passed as argument will reset the definition of the label.</p>				
RETURN VALUE	Upon successful completion, <code>setlabel()</code> returns 0; otherwise, it returns a non-zero value.				
EXAMPLES	<p>The following code (without previous call to <code>setlabel()</code>):</p> <pre>pfmt(stderr, MM_ERROR, "test:2:Cannot open file\n"); setlabel("UX:test"); pfmt(stderr, MM_ERROR, "test:2:Cannot open file\n");</pre> <p>will produce the following output:</p> <pre>ERROR: Cannot open file UX:test: ERROR: Cannot open file</pre>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<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>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>getopt(3C)</code> , <code>lfmt(3C)</code> , <code>pfmt(3C)</code> , <code>attributes(5)</code>				

setlocale(3C)

NAME	setlocale – modify and query a program’s locale																
SYNOPSIS	<pre>#include <locale.h> char *setlocale(int category, const char *locale);</pre>																
DESCRIPTION	<p>The <code>setlocale()</code> function selects the appropriate piece of the program’s locale as specified by the <i>category</i> and <i>locale</i> arguments. The <i>category</i> argument may have the following values: <code>LC_CTYPE</code>, <code>LC_NUMERIC</code>, <code>LC_TIME</code>, <code>LC_COLLATE</code>, <code>LC_MONETARY</code>, <code>LC_MESSAGES</code>, and <code>LC_ALL</code>. These names are defined in the <code><locale.h></code> header. The <code>LC_ALL</code> variable names all of a program’s locale categories.</p> <p>The <code>LC_CTYPE</code> variable affects the behavior of character handling functions such as <code>isdigit(3C)</code> and <code>tolower(3C)</code>, and multibyte character functions such as <code>mbtowc(3C)</code> and <code>wctomb(3C)</code>.</p> <p>The <code>LC_NUMERIC</code> variable affects the decimal point character and thousands separator character for the formatted input/output functions and string conversion functions.</p> <p>The <code>LC_TIME</code> variable affects the date and time format as delivered by <code>ascftime(3C)</code>, <code>cftime(3C)</code>, <code>getdate(3C)</code>, <code>strftime(3C)</code> and <code>strptime(3C)</code>.</p> <p>The <code>LC_COLLATE</code> variable affects the sort order produced by collating functions such as <code>strcoll(3C)</code> and <code>strxfrm(3C)</code>.</p> <p>The <code>LC_MONETARY</code> variable affects the monetary formatted information returned by <code>localeconv(3C)</code>.</p> <p>The <code>LC_MESSAGES</code> variable affects the behavior of messaging functions such as <code>dgettext(3C)</code>, <code>gettext(3C)</code>, and <code>gettext(3C)</code>.</p> <p>A value of "C" for <i>locale</i> specifies the traditional UNIX system behavior. At program startup, the equivalent of</p> <pre>setlocale(LC_ALL, "C")</pre> <p>is executed. This has the effect of initializing each category to the locale described by the environment "C".</p> <p>A value of "" for <i>locale</i> specifies that the locale should be taken from environment variables. The order in which the environment variables are checked for the various categories is given below:</p> <table border="1"><thead><tr><th>Category</th><th>1st Env Var</th><th>2nd Env Var</th><th>3rd Env Var</th></tr></thead><tbody><tr><td>LC_CTYPE:</td><td>LC_ALL</td><td>LC_CTYPE</td><td>LANG</td></tr><tr><td>LC_COLLATE:</td><td>LC_ALL</td><td>LC_COLLATE</td><td>LANG</td></tr><tr><td>LC_CTIME:</td><td>LC_ALL</td><td>LC_CTIME</td><td>LANG</td></tr></tbody></table>	Category	1st Env Var	2nd Env Var	3rd Env Var	LC_CTYPE:	LC_ALL	LC_CTYPE	LANG	LC_COLLATE:	LC_ALL	LC_COLLATE	LANG	LC_CTIME:	LC_ALL	LC_CTIME	LANG
Category	1st Env Var	2nd Env Var	3rd Env Var														
LC_CTYPE:	LC_ALL	LC_CTYPE	LANG														
LC_COLLATE:	LC_ALL	LC_COLLATE	LANG														
LC_CTIME:	LC_ALL	LC_CTIME	LANG														

Category	1st Env Var	2nd Env Var	3rd Env Var
LC_NUMERIC:	LC_ALL	LC_NUMERIC	LANG
LC_MONETARY:	LC_ALL	LC_MONETARY	LANG
LC_MESSAGES:	LC_ALL	LC_MESSAGES	LANG

If a pointer to a string is given for *locale*, `setlocale()` attempts to set the locale for the given category to *locale*. If `setlocale()` succeeds, *locale* is returned. If `setlocale()` fails, a null pointer is returned and the program's locale is not changed.

For category `LC_ALL`, the behavior is slightly different. If a pointer to a string is given for *locale* and `LC_ALL` is given for *category*, `setlocale()` attempts to set the locale for all the categories to *locale*. The *locale* may be a simple locale, consisting of a single locale, or a composite locale. If the locales for all the categories are the same after all the attempted locale changes, `setlocale()` will return a pointer to the common simple locale. If there is a mixture of locales among the categories, `setlocale()` will return a composite locale.

RETURN VALUES

Upon successful completion, `setlocale()` returns the string associated with the specified category for the new locale. Otherwise, `setlocale()` returns a null pointer and the program's locale is not changed.

A null pointer for *locale* causes `setlocale()` to return a pointer to the string associated with the *category* for the program's current locale. The program's locale is not changed.

The string returned by `setlocale()` is such that a subsequent call with that string and its associated *category* will restore that part of the program's locale. The string returned must not be modified by the program, but may be overwritten by a subsequent call to `setlocale()`.

ERRORS

No errors are defined.

FILES

`/usr/lib/locale/locale` locale database directory for *locale*

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

setlocale(3C)

SEE ALSO locale(1), ctype(3C), getdate(3C) gettext(3C), gettxt(3C), isdigit(3C), localeconv(3C), mbtowc(3C), strcoll(3C), strftime(3C), strptime(3C) strxfrm(3C) tolower(3C), wctomb(3C), libc(3LIB), attributes(5), environ(5), locale(5), standards(5)

NOTES It is unsafe for any thread to change locale (by calling `setlocale()` with a non-null locale argument) in a multithreaded application while any other thread in the application is using any locale-sensitive routine. To change locale in a multithreaded application, `setlocale()` should be called prior to using any locale-sensitive routine. Using `setlocale()` to query the current locale is safe and can be used anywhere in a multithreaded application except when some other thread is changing locale.

It is the user's responsibility to ensure that mixed locale categories are compatible. For example, setting `LC_CTYPE=C` and `LC_TIME=ja` (where `ja` indicates Japanese) will not work, because Japanese time cannot be represented in the "C" locale's ASCII codeset.

NAME	sigblock, sigmask, sigpause, sigsetmask – block signals
SYNOPSIS	<pre> /usr/ucb/cc [flag ...] file ... #include <signal.h> int sigblock (mask) ; int mask; int sigmask (signal) ; int signal; int sigpause (int mask) ; int mask; int sigsetmask (mask) ; int mask; </pre>
DESCRIPTION	<p>sigblock, sigmask, sigpause, sigsetmask – block signals</p> <p>sigblock () adds the signals specified in <i>mask</i> to the set of signals currently being blocked from delivery. Signals are blocked if the appropriate bit in <i>mask</i> is a 1; the macro sigmask is provided to construct the mask for a given <i>signal</i>. sigblock () returns the previous mask. The previous mask may be restored using sigsetmask ().</p> <p>sigpause () assigns <i>mask</i> to the set of masked signals and then waits for a signal to arrive; on return the set of masked signals is restored. <i>mask</i> is usually 0 to indicate that no signals are now to be blocked. sigpause () always terminates by being interrupted, returning -1 and setting errno to EINTR.</p> <p>sigsetmask () sets the current signal mask (those signals that are blocked from delivery). Signals are blocked if the corresponding bit in <i>mask</i> is a 1; the macro sigmask is provided to construct the mask for a given <i>signal</i>.</p> <p>In normal usage, a signal is blocked using sigblock (). To begin a critical section, variables modified on the occurrence of the signal are examined to determine that there is no work to be done, and the process pauses awaiting work by using sigpause () with the mask returned by sigblock ().</p> <p>It is not possible to block SIGKILL, SIGSTOP, or SIGCONT, this restriction is silently imposed by the system.</p>
RETURN VALUES	sigblock () and sigsetmask () return the previous set of masked signals. sigpause () returns -1 and sets errno to EINTR.
SEE ALSO	kill(2), sigaction(2), signal(3UCB), sigvec(3UCB)
NOTES	Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.

sigfpe(3C)

NAME	sigfpe – signal handling for specific SIGFPE codes
SYNOPSIS	<pre>#include <floatingpoint.h> #include <siginfo.h> sigfpe_handler_type sigfpe(sigfpe_code_type <i>code</i>, sigfpe_handler_type <i>hdl</i>);</pre>
DESCRIPTION	<p>The <code>sigfpe()</code> function allows signal handling to be specified for particular SIGFPE codes. A call to <code>sigfpe()</code> defines a new handler <i>hdl</i> for a particular SIGFPE <i>code</i> and returns the old handler as the value of the function <code>sigfpe()</code>. Normally handlers are specified as pointers to functions; the special cases SIGFPE_IGNORE, SIGFPE_ABORT, and SIGFPE_DEFAULT allow ignoring, dumping core using abort(3C), or default handling respectively. Default handling is to dump core using abort(3C).</p> <p>The <i>code</i> argument is usually one of the five IEEE 754-related SIGFPE codes:</p> <pre>FPE_FLTRES fp_inexact - floating-point inexact result FPE_FLTDIV fp_division - floating-point division by zero FPE_FLTUND fp_underflow - floating-point underflow FPE_FLTOVF fp_overflow - floating-point overflow FPE_FLTINV fp_invalid - floating-point invalid operation</pre> <p>Three steps are required to intercept an IEEE 754-related SIGFPE code with <code>sigfpe()</code>:</p> <ol style="list-style-type: none">1. Set up a handler with <code>sigfpe()</code>.2. Enable the relevant IEEE 754 trapping capability in the hardware, perhaps by using assembly-language instructions.3. Perform a floating-point operation that generates the intended IEEE 754 exception. <p>The <code>sigfpe()</code> function never changes floating-point hardware mode bits affecting IEEE 754 trapping. No IEEE 754-related SIGFPE signals will be generated unless those hardware mode bits are enabled.</p> <p>SIGFPE signals can be handled using <code>sigfpe()</code>, <code>sigaction(2)</code> or signal(3C). In a particular program, to avoid confusion, use only one of these interfaces to handle SIGFPE signals.</p>
EXAMPLES	<p>EXAMPLE 1 Example Of A User-Specified Signal Handler</p> <p>A user-specified signal handler might look like this:</p> <pre>#include <floatingpoint.h> #include <siginfo.h> #include <ucontext.h> /* * The sample_handler prints out a message then commits suicide. */ void sample_handler(int sig, siginfo_t *sip, ucontext_t *uap) { char *label; switch (sip->si_code) {</pre>

EXAMPLE 1 Example Of A User-Specified Signal Handler (Continued)

```

    case FPE_FLTINV: label = "invalid operand"; break;
    case FPE_FLTRES: label = "inexact"; break;
    case FPE_FLTDIV: label = "division-by-zero"; break;
    case FPE_FLTUND: label = "underflow"; break;
    case FPE_FLTOVF: label = "overflow"; break;
    default: label = "???"; break;
  }
  fprintf(stderr, "FP exception %s (0x%x) occurred at address %p.\n",
            label, sip->si_code, (void *) sip->si_addr);
  abort();
}

```

and it might be set up like this:

```

#include <floatingpoint.h>
#include <siginfo.h>
#include <ucontext.h>
extern void sample_handler(int, siginfo_t *, ucontext_t *);
main(void) {
    sigfpe_handler_type hdl, old_handler1, old_handler2;
    /*
     * save current fp_overflow and fp_invalid handlers; set the new
     * fp_overflow handler to sample_handler( ) and set the new
     * fp_invalid handler to SIGFPE_ABORT (abort on invalid)
     */
    hdl = (sigfpe_handler_type) sample_handler;
    old_handler1 = sigfpe(FPE_FLTOVF, hdl);
    old_handler2 = sigfpe(FPE_FLTINV, SIGFPE_ABORT);
    . . .
    /*
     * restore old fp_overflow and fp_invalid handlers
     */
    sigfpe(FPE_FLTOVF, old_handler1);
    sigfpe(FPE_FLTINV, old_handler2);
}

```

FILES /usr/include/floatingpoint.h

/usr/include/siginfo.h

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO sigaction(2), abort(3C), signal(3C), attributes(5), floatingpoint.h(3HEAD)

DIAGNOSTICS The sigfpe() function returns (void(*)()-1) if *code* is not zero or a defined SIGFPE code.

siginterrupt(3C)

NAME	siginterrupt – allow signals to interrupt functions						
SYNOPSIS	<pre>#include <signal.h> int siginterrupt(int <i>sig</i>, int <i>flag</i>);</pre>						
DESCRIPTION	<p>The <code>siginterrupt()</code> function changes the restart behavior when a function is interrupted by the specified signal. The function <code>siginterrupt(<i>sig</i>, <i>flag</i>)</code> has an effect as if implemented as:</p> <pre>siginterrupt(int sig, int flag) { int ret; struct sigaction act; (void) sigaction(sig, NULL, &act); if (flag) act.sa_flags &= SA_RESTART; else act.sa_flags = SA_RESTART; ret = sigaction(sig, &act, NULL); return ret; }</pre>						
RETURN VALUES	Upon successful completion, <code>siginterrupt()</code> returns 0. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.						
ERRORS	The <code>siginterrupt()</code> function will fail if: EINVAL The <i>sig</i> argument is not a valid signal number.						
USAGE	The <code>siginterrupt()</code> function supports programs written to historical system interfaces. A standard-conforming application, when being written or rewritten, should use <code>sigaction(2)</code> with the <code>SA_RESTART</code> flag instead of <code>siginterrupt()</code> .						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>sigaction(2)</code> , <code>signal.h(3HEAD)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	siginterrupt – allow signals to interrupt functions
SYNOPSIS	<pre>/usr/ucb/cc [flag ...] file ... int siginterrupt(sig, flag); int sig, flag;</pre>
DESCRIPTION	<p>siginterrupt () is used to change the function restart behavior when a function is interrupted by the specified signal. If the flag is false (0), then functions will be restarted if they are interrupted by the specified signal and no data has been transferred yet. System call restart is the default behavior when the signal(3C) routine is used.</p> <p>If the flag is true, (1), then restarting of functions is disabled. If a function is interrupted by the specified signal and no data has been transferred, the function will return -1 with <code>errno</code> set to <code>EINTR</code>. Interrupted functions that have started transferring data will return the amount of data actually transferred.</p> <p>Issuing a <code>siginterrupt ()</code> call during the execution of a signal handler will cause the new action to take place on the next signal to be caught.</p>
NOTES	<p>Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-threaded applications is unsupported.</p> <p>This library routine uses an extension of the sigvec(3UCB) function that is not available in 4.2 BSD, hence it should not be used if backward compatibility is needed.</p>
RETURN VALUES	A 0 value indicates that the call succeeded. A -1 value indicates that the call failed and <code>errno</code> is set to indicate the error.
ERRORS	<p>siginterrupt () may return the following error:</p> <p><code>EINVAL</code> <code>sig</code> is not a valid signal.</p>
SEE ALSO	sigblock(3UCB) , sigvec(3UCB) , signal(3C)

signal(3C)

NAME	signal, sigset, sighold, sigrelse, sigignore, sigpause – simplified signal management for application processes
SYNOPSIS	<pre>#include <signal.h> void (*signal (int sig, void (*disp) (int))) (int); void (*sigset (int sig, void (*disp) (int))) (int); int sighold (int sig); int sigrelse (int sig); int sigignore (int sig); int sigpause (int sig);</pre>
DESCRIPTION	<p>These functions provide simplified signal management for application processes. See <code>signal.h(3HEAD)</code> for an explanation of general signal concepts.</p> <p>The <code>signal()</code> and <code>sigset()</code> functions modify signal dispositions. The <i>sig</i> argument specifies the signal, which may be any signal except <code>SIGKILL</code> and <code>SIGSTOP</code>. The <i>disp</i> argument specifies the signal's disposition, which may be <code>SIG_DFL</code>, <code>SIG_IGN</code>, or the address of a signal handler. If <code>signal()</code> is used, <i>disp</i> is the address of a signal handler, and <i>sig</i> is not <code>SIGILL</code>, <code>SIGTRAP</code>, or <code>SIGPWR</code>, the system first sets the signal's disposition to <code>SIG_DFL</code> before executing the signal handler. If <code>sigset()</code> is used and <i>disp</i> is the address of a signal handler, the system adds <i>sig</i> to the calling process's signal mask before executing the signal handler; when the signal handler returns, the system restores the calling process's signal mask to its state prior to the delivery of the signal. In addition, if <code>sigset()</code> is used and <i>disp</i> is equal to <code>SIG_HOLD</code>, <i>sig</i> is added to the calling process's signal mask and the signal's disposition remains unchanged.</p> <p>The <code>sighold()</code> function adds <i>sig</i> to the calling process's signal mask.</p> <p>The <code>sigrelse()</code> function removes <i>sig</i> from the calling process's signal mask.</p> <p>The <code>sigignore()</code> function sets the disposition of <i>sig</i> to <code>SIG_IGN</code>.</p> <p>The <code>sigpause()</code> function removes <i>sig</i> from the calling process's signal mask and suspends the calling process until a signal is received.</p>
RETURN VALUES	<p>Upon successful completion, <code>signal()</code> returns the signal's previous disposition. Otherwise, it returns <code>SIG_ERR</code> and sets <code>errno</code> to indicate the error.</p> <p>Upon successful completion, <code>sigset()</code> returns <code>SIG_HOLD</code> if the signal had been blocked or the signal's previous disposition if it had not been blocked. Otherwise, it returns <code>SIG_ERR</code> and sets <code>errno</code> to indicate the error.</p> <p>Upon successful completion, <code>sighold()</code>, <code>sigrelse()</code>, <code>sigignore()</code>, and <code>sigpause()</code>, return 0. Otherwise, they return <code>-1</code> and set <code>errno</code> to indicate the error.</p>
ERRORS	These functions fail if:

EINTR A signal was caught during the execution `sigpause()`.

EINVAL The value of the `sig` argument is not a valid signal or is equal to `SIGKILL` or `SIGSTOP`.

USAGE The `sighold()` function used in conjunction with `sigrelse()` or `sigpause()` may be used to establish critical regions of code that require the delivery of a signal to be temporarily deferred.

If `signal()` or `sigset()` is used to set `SIGCHLD`'s disposition to a signal handler, `SIGCHLD` will not be sent when the calling process's children are stopped or continued.

If any of the above functions are used to set `SIGCHLD`'s disposition to `SIG_IGN`, the calling process's child processes will not create zombie processes when they terminate (see `exit(2)`). If the calling process subsequently waits for its children, it blocks until all of its children terminate; it then returns `-1` with `errno` set to `ECHILD` (see `wait(3C)` and `waitid(2)`).

The system guarantees that if more than one instance of the same signal is generated to a process, at least one signal will be received. It does not guarantee the reception of every generated signal.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `exit(2)`, `kill(2)`, `pause(2)`, `sigaction(2)`, `sigsend(2)`, `waitid(2)`, `signal.h(3HEAD)`, `wait(3C)`, `attributes(5)`, `standards(5)`

signal(3UCB)

NAME	signal – simplified software signal facilities
SYNOPSIS	<pre>/usr/ucb/cc [flag ...] file ... #include <signal.h> void (*signal(sig, func)) (); int sig; void (*func) ();</pre>
DESCRIPTION	<p>signal() is a simplified interface to the more general sigvec(3UCB) facility. Programs that use signal() in preference to sigvec() are more likely to be portable to all systems.</p> <p>A signal is generated by some abnormal event, initiated by a user at a terminal (quit, interrupt, stop), by a program error (bus error, etc.), by request of another program (kill), or when a process is stopped because it wishes to access its control terminal while in the background (see termio(7I)). Signals are optionally generated when a process resumes after being stopped, when the status of child processes changes, or when input is ready at the control terminal. Most signals cause termination of the receiving process if no action is taken; some signals instead cause the process receiving them to be stopped, or are simply discarded if the process has not requested otherwise. Except for the SIGKILL and SIGSTOP signals, the signal() call allows signals either to be ignored or to interrupt to a specified location. See sigvec(3UCB) for a complete list of the signals.</p> <p>If <i>func</i> is SIG_DFL, the default action for signal <i>sig</i> is reinstated; this default is termination (with a core image for starred signals) except for signals marked with • or a dagger. Signals marked with • are discarded if the action is SIG_DFL; signals marked with a dagger cause the process to stop. If <i>func</i> is SIG_IGN the signal is subsequently ignored and pending instances of the signal are discarded. Otherwise, when the signal occurs further occurrences of the signal are automatically blocked and <i>func</i> is called.</p> <p>A return from the function unblocks the handled signal and continues the process at the point it was interrupted.</p> <p>If a caught signal occurs during certain functions, terminating the call prematurely, the call is automatically restarted. In particular this can occur during a read(2) or write(2) on a slow device (such as a terminal; but not a file) and during a wait(3C).</p> <p>The value of signal() is the previous (or initial) value of <i>func</i> for the particular signal.</p> <p>After a fork(2) or vfork(2) the child inherits all signals. An exec(2) resets all caught signals to the default action; ignored signals remain ignored.</p>
RETURN VALUES	The previous action is returned on a successful call. Otherwise, -1 is returned and errno is set to indicate the error.
ERRORS	signal() will fail and no action will take place if the following occurs:

EINVAL *sig* is not a valid signal number, or is SIGKILL or SIGSTOP.

SEE ALSO kill(1), exec(2), fcntl(2), fork(2), getitimer(2), getrlimit(2), kill(2), read(2), sigaction(2), write(2), abort(3C), ptrace(3C), setjmp(3C), setjmp(3UCB), sigblock(3UCB), signal(3C), signal.h(3HEAD), sigstack(3UCB), sigvec(3UCB), wait(3C), wait(3UCB), termio(7I)

NOTES Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-threaded applications is unsupported.

The handler routine *func* can be declared:

```
void handler(signum) int signum;
```

Here *signum* is the signal number. See sigvec(3UCB) for more information.

sigsetops(3C)

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

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

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

SEE ALSO `sigaction(2)`, `sigpending(2)`, `sigprocmask(2)`, `sigsuspend(2)`, `signal.h(3HEAD)`, `attributes(5)`, `standards(5)`

sigstack(3C)

NAME	sigstack – set and/or get alternate signal stack context
SYNOPSIS	<pre>#include <signal.h> int sigstack(struct sigstack *ss, struct sigstack *oss);</pre>
DESCRIPTION	<p>The <code>sigstack()</code> function allows the calling process to indicate to the system an area of its address space to be used for processing signals received by the process.</p> <p>If the <code>ss</code> argument is not a null pointer, it must point to a <code>sigstack</code> structure. The length of the application-supplied stack must be at least <code>SIGSTKSZ</code> bytes. If the alternate signal stack overflows, the resulting behavior is undefined. (See <code>USAGE</code> below.)</p> <ul style="list-style-type: none">■ The value of the <code>ss_onstack</code> member indicates whether the process wants the system to use an alternate signal stack when delivering signals.■ The value of the <code>ss_sp</code> member indicates the desired location of the alternate signal stack area in the process' address space.■ If the <code>ss</code> argument is a null pointer, the current alternate signal stack context is not changed. <p>If the <code>oss</code> argument is not a null pointer, it points to a <code>sigstack</code> structure in which the current alternate signal stack context is placed. The value stored in the <code>ss_onstack</code> member of <code>oss</code> will be non-zero if the process is currently executing on the alternate signal stack. If the <code>oss</code> argument is a null pointer, the current alternate signal stack context is not returned.</p> <p>When a signal's action indicates its handler should execute on the alternate signal stack (specified by calling <code>sigaction(2)</code>), <code>sigstack()</code> checks to see if the process is currently executing on that stack. If the process is not currently executing on the alternate signal stack, the system arranges a switch to the alternate signal stack for the duration of the signal handler's execution.</p> <p>After a successful call to one of the <code>exec</code> functions, there are no alternate signal stacks in the new process image.</p>
RETURN VALUES	Upon successful completion, <code>sigstack()</code> returns 0. Otherwise, it returns -1 and sets <code>errno</code> to indicate the error.
ERRORS	The <code>sigstack()</code> function will fail if: EPERM An attempt was made to modify an active stack.
USAGE	A portable application, when being written or rewritten, should use <code>sigaltstack(2)</code> instead of <code>sigstack()</code> .

The direction of stack growth is not indicated in the historical definition of `struct sigstack`. The only way to portably establish a stack pointer is for the application to determine stack growth direction, or to allocate a block of storage and set the stack pointer to the middle. `sigstack()` may assume that the size of the signal stack is `SIGSTKSZ` as found in `<signal.h>`. An application that would like to specify a signal stack size other than `SIGSTKSZ` should use `sigaltstack(2)`.

Applications should not use `longjmp(3C)` to leave a signal handler that is running on a stack established with `sigstack()`. Doing so may disable future use of the signal stack. For abnormal exit from a signal handler, `siglongjmp(3C)`, `setcontext(2)`, or `swapcontext(3C)` may be used. These functions fully support switching from one stack to another.

The `sigstack()` function requires the application to have knowledge of the underlying system's stack architecture. For this reason, `sigaltstack(2)` is recommended over this function.

SEE ALSO `fork(2)`, `_longjmp(3C)`, `longjmp(3C)`, `setjmp(3C)`, `sigaltstack(2)`, `siglongjmp(3C)`, `sigsetjmp(3C)`

sigstack(3UCB)

NAME	sigstack – set and/or get signal stack context		
SYNOPSIS	<pre>/usr/ucb/cc [flag ...] file ... #include <signal.h> int sigstack(nss, oss); struct sigstack *nss, *oss;</pre>		
DESCRIPTION	<p>The <code>sigstack()</code> function allows users to define an alternate stack, called the “signal stack”, on which signals are to be processed. When a signal’s action indicates its handler should execute on the signal stack (specified with a <code>sigvec(3UCB)</code> call), the system checks to see if the process is currently executing on that stack. If the process is not currently executing on the signal stack, the system arranges a switch to the signal stack for the duration of the signal handler’s execution.</p> <p>A signal stack is specified by a <code>sigstack()</code> structure, which includes the following members:</p> <pre>char *ss_sp; /* signal stack pointer */ int ss_onstack; /* current status */</pre> <p>The <code>ss_sp</code> member is the initial value to be assigned to the stack pointer when the system switches the process to the signal stack. Note that, on machines where the stack grows downwards in memory, this is <i>not</i> the address of the beginning of the signal stack area. The <code>ss_onstack</code> member is zero or non-zero depending on whether the process is currently executing on the signal stack or not.</p> <p>If <code>nss</code> is not a null pointer, <code>sigstack()</code> sets the signal stack state to the value in the <code>sigstack()</code> structure pointed to by <code>nss</code>. If <code>nss</code> is a null pointer, the signal stack state will be unchanged. If <code>oss</code> is not a null pointer, the current signal stack state is stored in the <code>sigstack()</code> structure pointed to by <code>oss</code>.</p>		
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.		
ERRORS	The <code>sigstack()</code> function will fail and the signal stack context will remain unchanged if one of the following occurs.		
	<table><tr><td>EFAULT</td><td>Either <code>nss</code> or <code>oss</code> points to memory that is not a valid part of the process address space.</td></tr></table>	EFAULT	Either <code>nss</code> or <code>oss</code> points to memory that is not a valid part of the process address space.
EFAULT	Either <code>nss</code> or <code>oss</code> points to memory that is not a valid part of the process address space.		
SEE ALSO	<code>sigaltstack(2)</code> , <code>sigvec(3UCB)</code> , <code>signal(3C)</code>		
WARNINGS	Signal stacks are not “grown” automatically, as is done for the normal stack. If the stack overflows unpredictable results may occur.		
NOTES	Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-threaded applications is unsupported.		

NAME	sigvec – software signal facilities
SYNOPSIS	<pre> /usr/ucb/cc[flag ...] file... #include <signal.h> int sigvec (ssig, *nvec, *ovec); int sig; struct sigvec *nvec struct sigvec *ovec struct sigvec *nvec, *ovec; </pre>
DESCRIPTION	<p>The system defines a set of signals that may be delivered to a process. Signal delivery resembles the occurrence of a hardware interrupt: the signal is blocked from further occurrence, the current process context is saved, and a new one is built. A process may specify a <i>handler</i> to which a signal is delivered, or specify that a signal is to be <i>blocked</i> or <i>ignored</i>. A process may also specify that a default action is to be taken by the system when a signal occurs. Normally, signal handlers execute on the current stack of the process. This may be changed, on a per-handler basis, so that signals are taken on a special <i>signal stack</i>.</p> <p>All signals have the same <i>priority</i>. Signal routines execute with the signal that caused their invocation to be <i>blocked</i>, but other signals may yet occur. A global <i>signal mask</i> defines the set of signals currently blocked from delivery to a process. The signal mask for a process is initialized from that of its parent (normally 0). It may be changed with a <code>sigblock()</code> or <code>sigsetmask()</code> call, or when a signal is delivered to the process.</p> <p>A process may also specify a set of <i>flags</i> for a signal that affect the delivery of that signal.</p> <p>When a signal condition arises for a process, the signal is added to a set of signals pending for the process. If the signal is not currently <i>blocked</i> by the process then it is delivered to the process. When a signal is delivered, the current state of the process is saved, a new signal mask is calculated (as described below), and the signal handler is invoked. The call to the handler is arranged so that if the signal handling routine returns normally the process will resume execution in the context from before the signal's delivery. If the process wishes to resume in a different context, then it must arrange to restore the previous context itself.</p> <p>When a signal is delivered to a process a new signal mask is installed for the duration of the process' signal handler (or until a <code>sigblock()</code> or <code>sigsetmask()</code> call is made). This mask is formed by taking the current signal mask, adding the signal to be delivered, and ORing in the signal mask associated with the handler to be invoked.</p> <p>The action to be taken when the signal is delivered is specified by a <code>sigvec()</code> structure, which includes the following members:</p> <pre> void (*sv_handler)(); /* signal handler */ int sv_mask; /* signal mask to apply */ int sv_flags; /* see signal options */ </pre>

sigvec(3UCB)

```
#define SV_ONSTACK      /* take signal on signal stack */
#define SV_INTERRUPT   /* do not restart system on signal return */
#define SV_RESETHAND   /* reset handler to SIG_DFL when signal taken*/
```

If the `SV_ONSTACK` bit is set in the flags for that signal, the system will deliver the signal to the process on the signal stack specified with `sigstack(3UCB)` rather than delivering the signal on the current stack.

If `nvec` is not a NULL pointer, `sigvec()` assigns the handler specified by `sv_handler()`, the mask specified by `sv_mask()`, and the flags specified by `sv_flags()` to the specified signal. If `nvec` is a NULL pointer, `sigvec()` does not change the handler, mask, or flags for the specified signal.

The mask specified in `nvec` is not allowed to block `SIGKILL`, `SIGSTOP`, or `SIGCONT`. The system enforces this restriction silently.

If `ovect` is not a NULL pointer, the handler, mask, and flags in effect for the signal before the call to `sigvec()` are returned to the user. A call to `sigvec()` with `nvec` a NULL pointer and `ovect` not a NULL pointer can be used to determine the handling information currently in effect for a signal without changing that information.

The following is a list of all signals with names as in the include file `<signal.h>`:

<code>SIGHUP</code>	hangup
<code>SIGINT</code>	interrupt
<code>SIGQUIT*</code>	quit
<code>SIGILL*</code>	illegal instruction
<code>SIGTRAP*</code>	trace trap
<code>SIGABRT*</code>	abort (generated by <code>abort(3C)</code> routine)
<code>SIGEMT*</code>	emulator trap
<code>SIGFPE*</code>	arithmetic exception
<code>SIGKILL</code>	kill (cannot be caught, blocked, or ignored)
<code>SIGBUS*</code>	bus error
<code>SIGSEGV*</code>	segmentation violation
<code>SIGSYS*</code>	bad argument to function
<code>SIGPIPE</code>	write on a pipe or other socket with no one to read it
<code>SIGALRM</code>	alarm clock
<code>SIGTERM</code>	software termination signal
<code>SIGURG*</code>	urgent condition present on socket
<code>SIGSTOP**</code>	stop (cannot be caught, blocked, or ignored)

SIGTSTP**	stop signal generated from keyboard
SIGCONT*	continue after stop (cannot be blocked)
SIGCHLD*	child status has changed
SIGTTIN**	background read attempted from control terminal
SIGTTOU**	background write attempted to control terminal
SIGIO*	I/O is possible on a descriptor (see <code>fcntl(2)</code>)
SIGXCPU	cpu time limit exceeded (see <code>getrlimit(2)</code>)
SIGXFSZ	file size limit exceeded (see <code>getrlimit(2)</code>)
SIGVTALRM	virtual time alarm; see <code>setitimer()</code> on <code>getitimer(2)</code>
SIGPROF	profiling timer alarm; see <code>setitimer()</code> on <code>getitimer(2)</code>
SIGWINCH*	window changed (see <code>termio(7I)</code>)
SIGLOST	resource lost (see <code>lockd(1M)</code>)
SIGUSR1	user-defined signal 1
SIGUSR2	user-defined signal 2

The starred signals in the list above cause a core image if not caught or ignored.

Once a signal handler is installed, it remains installed until another `sigvec()` call is made, or an `execve(2)` is performed, unless the `SV_RESETHAND` bit is set in the flags for that signal. In that case, the value of the handler for the caught signal will be set to `SIG_DFL` before entering the signal-catching function, unless the signal is `SIGILL`, `SIGPWR`, or `SIGTRAP`. Also, if this bit is set, the bit for that signal in the signal mask will not be set; unless the signal mask associated with that signal blocks that signal, further occurrences of that signal will not be blocked. The `SV_RESETHAND` flag is not available in 4.2BSD, hence it should not be used if backward compatibility is needed.

The default action for a signal may be reinstated by setting the signal's handler to `SIG_DFL`; this default is termination except for signals marked with * or **. Signals marked with * are discarded if the action is `SIG_DFL`; signals marked with ** cause the process to stop. If the process is terminated, a "core image" will be made in the current working directory of the receiving process if the signal is one for which an asterisk appears in the above list (see `core(4)`).

If the handler for that signal is `SIG_IGN`, the signal is subsequently ignored, and pending instances of the signal are discarded.

sigvec(3UCB)

If a caught signal occurs during certain functions, the call is normally restarted. The call can be forced to terminate prematurely with an `EINTR` error return by setting the `SV_INTERRUPT` bit in the flags for that signal. The `SV_INTERRUPT` flag is not available in 4.2BSD, hence it should not be used if backward compatibility is needed. The affected functions are `read(2)` or `write(2)` on a slow device (such as a terminal or pipe or other socket, but not a file) and during a `wait(3C)`.

After a `fork(2)` or `vfork(2)` the child inherits all signals, the signal mask, the signal stack, and the restart/interrupt and reset-signal-handler flags.

The `execve(2)` call resets all caught signals to default action and resets all signals to be caught on the user stack. Ignored signals remain ignored; the signal mask remains the same; signals that interrupt functions continue to do so.

The accuracy of *addr* is machine dependent. For example, certain machines may supply an address that is on the same page as the address that caused the fault. If an appropriate *addr* cannot be computed it will be set to `SIG_NOADDR`.

RETURN VALUES

A 0 value indicates that the call succeeded. A -1 return value indicates that an error occurred and `errno` is set to indicate the reason.

ERRORS

`sigvec()` will fail and no new signal handler will be installed if one of the following occurs:

`EFAULT` Either *nvec* or *ovec* is not a `NULL` pointer and points to memory that is not a valid part of the process address space.

`EINVAL` *sig* is not a valid signal number, or, `SIGKILL`, or `SIGSTOP`.

SEE ALSO

`intro(2)`, `exec(2)`, `fcntl(2)`, `fork(2)`, `getitimer(2)`, `getrlimit(2)`, `ioctl(2)`, `kill(2)`, `read(2)`, `umask(2)`, `vfork(2)`, `write(2)`, `ptrace(3C)`, `setjmp(3C)`, `sigblock(3UCB)`, `signal(3C)`, `signal(3UCB)`, `sigstack(3UCB)`, `wait(3C)`, `wait(3UCB)`, `core(4)`, `streamio(7I)`, `termio(7I)`

NOTES

Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.

`SIGPOLL` is a synonym for `SIGIO`. A `SIGIO` will be issued when a file descriptor corresponding to a `STREAMS` (see `intro(2)`) file has a "selectable" event pending. Unless that descriptor has been put into asynchronous mode (see `fcntl(2)`), a process may specifically request that this signal be sent using the `I_SETSIG` `ioctl(2)` call (see `streamio(7I)`). Otherwise, the process will never receive `SIGPOLLs0`.

The handler routine can be declared:

```
void handler(int sig, int code, struct sigcontext *scp, char *addr);
```

sigvec(3UCB)

Here *sig* is the signal number; *code* is a parameter of certain signals that provides additional detail; *scp* is a pointer to the `sigcontext` structure (defined in `signal.h`), used to restore the context from before the signal; and *addr* is additional address information.

The signals `SIGKILL`, `SIGSTOP`, and `SIGCONT` cannot be ignored.

sleep(3C)

NAME	sleep – suspend execution for an interval of time						
SYNOPSIS	<pre>#include <unistd.h> unsigned int sleep(unsigned int <i>seconds</i>);</pre>						
DESCRIPTION	<p>The current process is suspended from execution for the number of <i>seconds</i> specified by the argument. The actual suspension time may be less than that requested because any caught signal will terminate the <code>sleep()</code> following execution of that signal's catching routine. Also, the suspension time may be longer than requested by an arbitrary amount because of the scheduling of other activity in the system. The value returned by <code>sleep()</code> will be the "unslept" amount (the requested time minus the time actually slept) in case the caller had an alarm set to go off earlier than the end of the requested <code>sleep()</code> time, or premature arousal because of another caught signal.</p> <p>In a single-threaded program (one not linked with <code>-lthread</code> or <code>-lpthread</code>), the routine is implemented by setting an alarm signal and pausing until it (or some other signal) occurs. The previous state of the alarm signal is saved and restored. The calling program may have set up an alarm signal before calling <code>sleep()</code>. If the <code>sleep()</code> time exceeds the time until such alarm signal, the process sleeps only until the alarm signal would have occurred. The caller's alarm catch routine is executed just before the <code>sleep()</code> routine returns. But if the <code>sleep()</code> time is less than the time till such alarm, the prior alarm time is reset to go off at the same time it would have without the intervening <code>sleep()</code>.</p> <p>In a multithreaded program (one linked with <code>-lthread</code> or <code>-lpthread</code>), the routine is implemented with a call to the <code>nanosleep(3RT)</code> function and does not modify the state of the alarm signal.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>Async-Signal-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Async-Signal-Safe						
SEE ALSO	<code>alarm(2)</code> , <code>pause(2)</code> , <code>signal(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						
NOTES	<p>In a single-threaded program, the <code>SIGALRM</code> signal should not be blocked or ignored during a call to <code>sleep()</code>. This restriction does not apply to a multithreaded program.</p> <p>In a multithreaded program, only the invoking thread is suspended from execution.</p>						

NAME	sleep – suspend execution for interval				
SYNOPSIS	<pre>/usr/ucb/cc [flag ...] file ...</pre> <pre>int sleep(seconds);</pre> <pre>unsigned seconds;</pre>				
DESCRIPTION	<p><code>sleep()</code> suspends the current process from execution for the number of seconds specified by the argument. The actual suspension time may be up to 1 second less than that requested, because scheduled wakeups occur at fixed 1-second intervals, and may be an arbitrary amount longer because of other activity in the system.</p> <p><code>sleep()</code> is implemented by setting an interval timer and pausing until it expires. The previous state of this timer is saved and restored. If the sleep time exceeds the time to the expiration of the previous value of the timer, the process sleeps only until the timer would have expired, and the signal which occurs with the expiration of the timer is sent one second later.</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <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	<code>alarm(2)</code> , <code>getitimer(2)</code> , <code>longjmp(3C)</code> , <code>siglongjmp(3C)</code> , <code>sleep(3C)</code> , <code>usleep(3C)</code> , <code>attributes(5)</code>				
NOTES	<p>Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.</p> <p>SIGALRM should <i>not</i> be blocked or ignored during a call to <code>sleep()</code>. Only a prior call to <code>alarm(2)</code> should generate SIGALRM for the calling process during a call to <code>sleep()</code>. A signal-catching function should <i>not</i> interrupt a call to <code>sleep()</code> to call <code>siglongjmp(3C)</code> or <code>longjmp(3C)</code> to restore an environment saved prior to the <code>sleep()</code> call.</p>				
WARNINGS	<code>sleep()</code> is slightly incompatible with <code>alarm(2)</code> . Programs that do not execute for at least one second of clock time between successive calls to <code>sleep()</code> indefinitely delay the alarm signal. Use <code>sleep(3C)</code> . Each <code>sleep(3C)</code> call postpones the alarm signal that would have been sent during the requested sleep period to occur one second later.				

ssignal(3C)

NAME	ssignal, gsignal – software signals				
SYNOPSIS	<pre>#include <signal.h> void(*ssignal (int sig, int (*action) (int))) (int); int gsignal (int sig);</pre>				
DESCRIPTION	<p>The <code>ssignal()</code> and <code>gsignal()</code> functions implement a software facility similar to signal(3C). This facility is made available to users for their own purposes.</p>				
ssignal()	<p>Software signals made available to users are associated with integers in the inclusive range 1 through 17. A call to <code>ssignal()</code> associates a procedure, <i>action</i>, with the software signal <i>sig</i>; the software signal, <i>sig</i>, is raised by a call to <code>gsignal()</code>. Raising a software signal causes the action established for that signal to be taken.</p> <p>The first argument to <code>ssignal()</code> is a number identifying the type of signal for which an action is to be established. The second argument defines the action; it is either the name of a (user-defined) <i>action function</i> or one of the manifest constants <code>SIG_DFL</code> (default) or <code>SIG_IGN</code> (ignore). The <code>ssignal()</code> function returns the action previously established for that signal type; if no action has been established or the signal number is illegal, <code>ssignal()</code> returns <code>SIG_DFL</code>.</p>				
gsignal()	<p>The <code>gsignal()</code> raises the signal identified by its argument, <i>sig</i>.</p> <p>If an action function has been established for <i>sig</i>, then that action is reset to <code>SIG_DFL</code> and the action function is entered with argument <i>sig</i>. The <code>gsignal()</code> function returns the value returned to it by the action function.</p> <p>If the action for <i>sig</i> is <code>SIG_IGN</code>, <code>gsignal()</code> returns the value 1 and takes no other action.</p> <p>If the action for <i>sig</i> is <code>SIG_DFL</code>, <code>gsignal()</code> returns the value 0 and takes no other action.</p> <p>If <i>sig</i> has an illegal value or no action was ever specified for <i>sig</i>, <code>gsignal()</code> returns the value 0 and takes no other action.</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>Unsafe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Unsafe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Unsafe				
SEE ALSO	raise(3C) , signal(3C) , attributes(5)				

NAME	stack_getbounds – retrieve stack boundaries						
SYNOPSIS	<pre>#include <ucontext.h> int stack_getbounds(stack_t *sp);</pre>						
DESCRIPTION	<p>The <code>stack_getbounds()</code> function retrieves the stack boundaries that the calling thread is currently operating on. If the thread is currently operating on the alternate signal stack, this function will retrieve the bounds of that stack.</p> <p>If successful, <code>stack_getbounds()</code> sets the <code>ss_sp</code> member of the <code>stack_t</code> structure pointed to by <code>sp</code> to the base of the stack region and the <code>ss_size</code> member to its size (maximum extent) in bytes. The <code>ss_flags</code> member is set to <code>SS_ONSTACK</code> if the calling thread is executing on its alternate signal stack, and zero otherwise.</p>						
RETURN VALUES	Upon successful completion, <code>stack_getbounds()</code> returns 0. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.						
ERRORS	<p>The <code>stack_getbounds()</code> function will fail if:</p> <p><code>EFAULT</code> The <code>sp</code> argument does not refer to a valid address.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Evolving</td> </tr> <tr> <td>MT-Level</td> <td>Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
MT-Level	Async-Signal-Safe						
SEE ALSO	<code>getustack(2)</code> , <code>sigaction(2)</code> , <code>sigaltstack(2)</code> , <code>stack_setbounds(3C)</code> , <code>attributes(5)</code>						

`_stack_grow(3C)`

NAME	<code>_stack_grow</code> – express an intention to extend the stack						
SYNOPSIS	<pre>#include <ucontext.h> void * _stack_grow(void *<i>addr</i>);</pre>						
DESCRIPTION	<p>The <code>_stack_grow()</code> function indicates to the system that the stack is about to be extended to the address specified by <i>addr</i>. If extending the stack to this address would violate the stack boundaries as retrieved by <code>stack_getbounds(3C)</code>, a SIGSEGV is raised.</p> <p>If the disposition of SIGSEGV is SIG_DFL, the process is terminated and a core dump is generated. If the application has installed its own SIGSEGV handler to run on the alternate signal stack, the signal information passed to the handler will be such that a call to <code>stack_violation(3C)</code> with these parameters returns 1.</p> <p>The <i>addr</i> argument is a biased stack pointer value. See the Solaris 64-bit Developer's Guide.</p> <p>This function has no effect if the specified address, <i>addr</i>, is within the bounds of the current stack.</p>						
RETURN VALUES	If the <code>_stack_grow()</code> function succeeds and does not detect a stack violation, it returns <i>addr</i> .						
ERRORS	No errors are defined.						
USAGE	<p>The <code>_stack_grow()</code> function does not actually adjust the stack pointer register. The caller is responsible for manipulating the stack pointer register once <code>_stack_grow()</code> returns.</p> <p>The <code>_stack_grow()</code> function is typically invoked by code created by the compilation environment prior to executing code that modifies the stack pointer. It can also be used by hand-written assembly routines to allocate stack-based storage safely.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Evolving</td></tr><tr><td>MT-Level</td><td>Async-Signal-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
MT-Level	Async-Signal-Safe						
SEE ALSO	<code>stack_getbounds(3C)</code> , <code>stack_inbounds(3C)</code> , <code>stack_violation(3C)</code> , <code>attributes(5)</code>						
	Solaris 64-bit Developer's Guide						

NAME	stack_inbounds – determine if address is within stack boundaries
SYNOPSIS	<pre>#include <ucontext.h> int stack_inbounds(void *addr);</pre>
DESCRIPTION	The <code>stack_inbounds()</code> function returns a boolean value indicating whether the address specified by <i>addr</i> is within the boundaries of the stack of the calling thread. The address is compared to the stack boundary information returned by a call to stack_getbounds(3C) .
RETURN VALUES	The <code>stack_inbounds()</code> function returns 0 to indicate that <i>addr</i> is not within the current stack bounds, or a non-zero value to indicate that <i>addr</i> is within the stack bounds.
ERRORS	No errors are defined.
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

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

SEE ALSO [stack_getbounds\(3C\)](#), [attributes\(5\)](#)

stack_setbounds(3C)

NAME	stack_setbounds – update stack boundaries						
SYNOPSIS	<pre>#include <ucontext.h> int stack_setbounds(const stack_t *sp);</pre>						
DESCRIPTION	The <code>stack_setbounds()</code> function updates the current base and bounds of the stack for the current thread to the bounds specified by the <code>stack_t</code> structure pointed to by <code>sp</code> . The <code>ss_sp</code> member refers to the virtual address of the base of the stack memory. The <code>ss_size</code> member refers to the size of the stack in bytes. The <code>ss_flags</code> member must be set to 0.						
RETURN VALUES	Upon successful completion, <code>stack_setbounds()</code> returns 0. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.						
ERRORS	The <code>stack_setbounds()</code> function will fail if: EFAULT The <code>sp</code> argument does not refer to a valid address or the <code>ss_sp</code> member of the <code>stack_t</code> structure pointed to by <code>sp</code> points to an illegal address. EINVAL The <code>ss_sp</code> member of the <code>stack_t</code> structure pointed to by <code>sp</code> is not properly aligned, the <code>ss_size</code> member is too small or is not properly aligned, or the <code>ss_flags</code> member is non-zero.						
USAGE	The <code>stack_setbounds()</code> function is intended for use by applications that are managing their own alternate stacks.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Evolving</td></tr><tr><td>MT-Level</td><td>Async-Signal-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Evolving	MT-Level	Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Evolving						
MT-Level	Async-Signal-Safe						
SEE ALSO	<code>getustack(2)</code> , <code>_stack_grow(3C)</code> , <code>stack_getbounds(3C)</code> , <code>stack_inbounds(3C)</code> , <code>stack_violation(3C)</code> , <code>attributes(5)</code>						

NAME	stack_violation – determine stack boundary violation event
SYNOPSIS	<pre>#include <ucontext.h> int stack_violation(int sig, const siginfo_t *sip, const ucontext_t *ucp);</pre>
DESCRIPTION	The stack_violation() function returns a boolean value indicating whether the signal, <i>sig</i> , and accompanying signal information, <i>sip</i> , and saved context, <i>ucp</i> , represent a stack boundary violation event or a stack overflow.
RETURN VALUES	The stack_violation() function returns 0 if the signal does not represent a stack boundary violation event and 1 if the signal does represent a stack boundary violation event.
ERRORS	No errors are defined.
EXAMPLES	<p>EXAMPLE 1 Set up a signal handler to run on an alternate stack.</p> <p>The following example sets up a signal handler for SIGSEGV to run on an alternate signal stack. For each signal it handles, the handler emits a message to indicate if the signal was produced due to a stack boundary violation.</p> <pre>#include <stdlib.h> #include <unistd.h> #include <ucontext.h> #include <signal.h> static void handler(int sig, siginfo_t *sip, void *p) { ucontext_t *ucp = p; const char *str; if (stack_violation(sig, sip, ucp)) str = "stack violation.\n"; else str = "no stack violation.\n"; (void) write(STDERR_FILENO, str, strlen(str)); exit(1); } int main(int argc, char **argv) { struct sigaction sa; stack_t altstack; altstack.ss_size = SIGSTKSZ; altstack.ss_sp = malloc(SIGSTKSZ); altstack.ss_flags = 0; (void) sigaltstack(&altstack, NULL);</pre>

stack_violation(3C)

EXAMPLE 1 Set up a signal handler to run on an alternate stack. *(Continued)*

```
sa.sa_sigaction = handler;
(void) sigfillset(&sa.sa_mask);
sa.sa_flags = SA_ONSTACK | SA_SIGINFO;
(void) sigaction(SIGSEGV, &sa, NULL);

/*
 * The application is now set up to use stack_violation(3C).
 */

return (0);
}
```

USAGE An application typically uses `stack_violation()` in a signal handler that has been installed for `SIGSEGV` using `sigaction(2)` with the `SA_SIGINFO` flag set and is configured to run on an alternate signal stack.

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

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

SEE ALSO `sigaction(2)`, `sigaltstack(2)`, `stack_getbounds(3C)`, `stack_inbounds(3C)`, `stack_setbounds(3C)`, `attributes(5)`

NAME	stdio – standard buffered input/output package																		
SYNOPSIS	<pre>#include <stdio.h> extern FILE *stdin; extern FILE *stdout; extern FILE *stderr;</pre>																		
DESCRIPTION	<p>The functions described in the entries of section 3S of this manual constitute an efficient, user-level I/O buffering scheme. The in-line macros <code>getc()</code> and <code>putc()</code> handle characters quickly. The macros <code>getchar(3C)</code> and <code>putchar(3C)</code>, and the higher-level routines <code>fgetc(3C)</code>, <code>fgets(3C)</code>, <code>fprintf(3C)</code>, <code>fputc(3C)</code>, <code>fputs(3C)</code>, <code>fread(3C)</code>, <code>fscanf(3C)</code>, <code>fwrite(3C)</code>, <code>gets(3C)</code>, <code>getw(3C)</code>, <code>printf(3C)</code>, <code>puts(3C)</code>, <code>putw(3C)</code>, and <code>scanf(3C)</code> all use or act as if they use <code>getc()</code> and <code>putc()</code>; they can be freely intermixed.</p> <p>A file with associated buffering is called a <i>stream</i> (see <code>intro(3)</code>) and is declared to be a pointer to a defined type <code>FILE</code>. The <code>fopen(3C)</code> function creates certain descriptive data for a stream and returns a pointer to designate the stream in all further transactions. Normally, there are three open streams with constant pointers declared in the <code><stdio.h></code> header and associated with the standard open files:</p> <table border="0"> <tr> <td><code>stdin</code></td> <td>standard input file</td> </tr> <tr> <td><code>stdout</code></td> <td>standard output file</td> </tr> <tr> <td><code>stderr</code></td> <td>standard error file</td> </tr> </table> <p>The following symbolic values in <code><unistd.h></code> define the file descriptors that will be associated with the C-language <code>stdin</code>, <code>stdout</code> and <code>stderr</code> when the application is started:</p> <table border="0"> <tr> <td><code>STDIN_FILENO</code></td> <td>Standard input value</td> <td>0</td> <td><code>stdin</code></td> </tr> <tr> <td><code>STDOUT_FILENO</code></td> <td>Standard output value</td> <td>1</td> <td><code>stdout</code></td> </tr> <tr> <td><code>STDERR_FILENO</code></td> <td>Standard error value</td> <td>2</td> <td><code>stderr</code></td> </tr> </table> <p>The constant <code>NULL</code> designates a null pointer.</p> <p>The integer-constant <code>EOF</code> is returned upon end-of-file or error by most integer functions that deal with streams (see the individual descriptions for details).</p> <p>The integer constant <code>BUFSIZ</code> specifies the size of the buffers used by the particular implementation.</p> <p>The integer constant <code>FILENAME_MAX</code> specifies the number of bytes needed to hold the longest pathname of a file allowed by the implementation. If the system does not impose a maximum limit, this value is the recommended size for a buffer intended to hold a file's pathname.</p>	<code>stdin</code>	standard input file	<code>stdout</code>	standard output file	<code>stderr</code>	standard error file	<code>STDIN_FILENO</code>	Standard input value	0	<code>stdin</code>	<code>STDOUT_FILENO</code>	Standard output value	1	<code>stdout</code>	<code>STDERR_FILENO</code>	Standard error value	2	<code>stderr</code>
<code>stdin</code>	standard input file																		
<code>stdout</code>	standard output file																		
<code>stderr</code>	standard error file																		
<code>STDIN_FILENO</code>	Standard input value	0	<code>stdin</code>																
<code>STDOUT_FILENO</code>	Standard output value	1	<code>stdout</code>																
<code>STDERR_FILENO</code>	Standard error value	2	<code>stderr</code>																

stdio(3C)

The integer constant `FOPEN_MAX` specifies the minimum number of files that the implementation guarantees can be open simultaneously. Note that no more than 255 files may be opened using `fopen()`, and only file descriptors 0 through 255 can be used in a stream.

The functions and constants mentioned in the entries of section 3S of this manual are declared in that header and need no further declaration. The constants and the following “functions” are implemented as macros (redeclaration of these names is perilous): `getc()`, `getchar()`, `putc()`, `putchar()`, `ferror(3C)`, `feof(3C)`, `clearerr(3C)`, and `fileno(3C)`. There are also function versions of `getc()`, `getchar()`, `putc()`, `putchar()`, `ferror()`, `feof()`, `clearerr()`, and `fileno()`.

Output streams, with the exception of the standard error stream `stderr`, are by default buffered if the output refers to a file and line-buffered if the output refers to a terminal. The standard error output stream `stderr` is by default unbuffered, but use of `freopen()` (see `fopen(3C)`) will cause it to become buffered or line-buffered. When an output stream is unbuffered, information is queued for writing on the destination file or terminal as soon as written; when it is buffered, many characters are saved up and written as a block. When it is line-buffered, each line of output is queued for writing on the destination terminal as soon as the line is completed (that is, as soon as a new-line character is written or terminal input is requested). The `setbuf()` or `setvbuf()` functions (both described on the `setbuf(3C)` manual page) may be used to change the stream’s buffering strategy.

Interactions of Other FILE-Type C Functions

A single open file description can be accessed both through streams and through file descriptors. Either a file descriptor or a stream will be called a *handle* on the open file description to which it refers; an open file description may have several handles.

Handles can be created or destroyed by user action without affecting the underlying open file description. Some of the ways to create them include `fcntl(2)`, `dup(2)`, `fdopen(3C)`, `fileno(3C)` and `fork(2)` (which duplicates existing ones into new processes). They can be destroyed by at least `fclose(3C)` and `close(2)`, and by the `exec` functions (see `exec(2)`), which close some file descriptors and destroy streams.

A file descriptor that is never used in an operation and could affect the file offset (for example `read(2)`, `write(2)`, or `lseek(2)`) is not considered a handle in this discussion, but could give rise to one (as a consequence of `fdopen()`, `dup()`, or `fork()`, for example). This exception does include the file descriptor underlying a stream, whether created with `fopen()` or `fdopen()`, as long as it is not used directly by the application to affect the file offset. (The `read()` and `write()` functions implicitly affect the file offset; `lseek()` explicitly affects it.)

If two or more handles are used, and any one of them is a stream, their actions shall be coordinated as described below. If this is not done, the result is undefined.

A handle that is a stream is considered to be closed when either an `fclose()` or `freopen(3C)` is executed on it (the result of `freopen()` is a new stream for this discussion, which cannot be a handle on the same open file description as its previous

value) or when the process owning that stream terminates the `exit(2)` or `abort(3C)`. A file descriptor is closed by `close()`, `_exit()` (see `exit(2)`), or by one of the `exec` functions when `FD_CLOEXEC` is set on that file descriptor.

For a handle to become the active handle, the actions below must be performed between the last other user of the first handle (the current active handle) and the first other user of the second handle (the future active handle). The second handle then becomes the active handle. All activity by the application affecting the file offset on the first handle shall be suspended until it again becomes the active handle. (If a stream function has as an underlying function that affects the file offset, the stream function will be considered to affect the file offset. The underlying functions are described below.)

The handles need not be in the same process for these rules to apply. Note that after a `fork()`, two handles exist where one existed before. The application shall assure that, if both handles will ever be accessed, that they will both be in a state where the other could become the active handle first. The application shall prepare for a `fork()` exactly as if it were a change of active handle. (If the only action performed by one of the processes is one of the `exec` functions or `_exit()`, the handle is never accessed in that process.)

1. For the first handle, the first applicable condition below shall apply. After the actions required below are taken, the handle may be closed if it is still open.
 - a. If it is a file descriptor, no action is required.
 - b. If the only further action to be performed on any handle to this open file description is to close it, no action need be taken.
 - c. If it is a stream that is unbuffered, no action need be taken.
 - d. If it is a stream that is line-buffered and the last character written to the stream was a newline (that is, as if a `putc('\n')` was the most recent operation on that stream), no action need be taken.
 - e. If it is a stream that is open for writing or append (but not also open for reading), either an `fflush(3C)` shall occur or the stream shall be closed.
 - f. If the stream is open for reading and it is at the end of the file (`feof(3C)` is true), no action need be taken.
 - g. If the stream is open with a mode that allows reading and the underlying open file description refers to a device that is capable of seeking, either an `fflush()` shall occur or the stream shall be closed.
 - h. Otherwise, the result is undefined.
2. For the second handle: if any previous active handle has called a function that explicitly changed the file offset, except as required above for the first handle, the application shall perform an `lseek()` or an `fseek(3C)` (as appropriate to the type of the handle) to an appropriate location.
3. If the active handle ceases to be accessible before the requirements on the first handle above have been met, the state of the open file description becomes undefined. This might occur, for example, during a `fork()` or an `_exit()`.

stdio(3C)

4. The `exec` functions shall be considered to make inaccessible all streams that are open at the time they are called, independent of what streams or file descriptors may be available to the new process image.
5. Implementation shall assure that an application, even one consisting of several processes, shall yield correct results (no data is lost or duplicated when writing, all data is written in order, except as requested by seeks) when the rules above are followed, regardless of the sequence of handles used. If the rules above are not followed, the result is unspecified. When these rules are followed, it is implementation defined whether, and under what conditions, all input is seen exactly once.

Use of stdio in Multithreaded Applications

All the `stdio` functions are safe unless they have the `_unlocked` suffix. Each `FILE` pointer has its own lock to guarantee that only one thread can access it. In the case that output needs to be synchronized, the lock for the `FILE` pointer can be acquired before performing a series of `stdio` operations. For example:

```
FILE iop;
flockfile(iop);
fprintf(iop, "hello ");
fprintf(iop, "world");
fputc(iop, 'a');
funlockfile(iop);
```

will print everything out together, blocking other threads that might want to write to the same file between calls to `fprintf()`.

An unlocked interface is available in case performance is an issue. For example:

```
flockfile(iop);
while (!feof(iop)) {
    *c++ = getc_unlocked(iop);
}
funlockfile(iop);
```

RETURN VALUES

Invalid stream pointers usually cause grave disorder, possibly including program termination. Individual function descriptions describe the possible error conditions.

SEE ALSO

`close(2)`, `lseek(2)`, `open(2)`, `pipe(2)`, `read(2)`, `write(2)`, `ctermid(3C)`, `cuserid(3C)`, `fclose(3C)`, `ferror(3C)`, `fopen(3C)`, `fread(3C)`, `fseek(3C)`, `flockfile(3C)`, `getc(3C)`, `gets(3C)`, `popen(3C)`, `printf(3C)`, `putc(3C)`, `puts(3C)`, `scanf(3C)`, `setbuf(3C)`, `system(3C)`, `tmpfile(3C)`, `tmpnam(3C)`, `ungetc(3C)`

NAME	str2sig, sig2str – translation between signal name and signal number
SYNOPSIS	<pre>#include <signal.h> int str2sig(const char *str, int *signum); int sig2str(int signum, char *str);</pre>
DESCRIPTION	<p>The <code>str2sig()</code> function translates the signal name <i>str</i> to a signal number, and stores that result in the location referenced by <i>signum</i>. The name in <i>str</i> can be either the symbol for that signal, without the "SIG" prefix, or a decimal number. All the signal symbols defined in <code><sys/signal.h></code> are recognized. This means that both "CLD" and "CHLD" are recognized and return the same signal number, as do both "POLL" and "IO". For access to the signals in the range SIGRTMIN to SIGRTMAX, the first four signals match the strings "RTMIN", "RTMIN+1", "RTMIN+2", and "RTMIN+3" and the last four match the strings "RTMAX-3", "RTMAX-2", "RTMAX-1", and "RTMAX".</p> <p>The <code>sig2str()</code> function translates the signal number <i>signum</i> to the symbol for that signal, without the "SIG" prefix, and stores that symbol at the location specified by <i>str</i>. The storage referenced by <i>str</i> should be large enough to hold the symbol and a terminating null byte. The symbol SIG2STR_MAX defined by <code><signal.h></code> gives the maximum size in bytes required.</p>
RETURN VALUES	<p>The <code>str2sig()</code> function returns 0 if it recognizes the signal name specified in <i>str</i>; otherwise, it returns -1.</p> <p>The <code>sig2str()</code> function returns 0 if the value <i>signum</i> corresponds to a valid signal number; otherwise, it returns -1.</p>
EXAMPLES	<p>EXAMPLE 1 A sample program using the <code>str2sig()</code> function.</p> <pre>int i; char buf[SIG2STR_MAX]; /*storage for symbol */ str2sig("KILL",&i); /*stores 9 in i */ str2sig("9", &i); /* stores 9 in i */ sig2str(SIGKILL,buf); /* stores "KILL" in buf */ sig2str(9,buf); /* stores "KILL" in buf */</pre>
SEE ALSO	kill(1), strsignal(3C)

strcoll(3C)

NAME	strcoll – string collation								
SYNOPSIS	<pre>#include <string.h> int strcoll(const char *s1, const char *s2);</pre>								
DESCRIPTION	<p>Both <code>strcoll()</code> and strxfrm(3C) provide for locale-specific string sorting. <code>strcoll()</code> is intended for applications in which the number of comparisons per string is small. When strings are to be compared a number of times, strxfrm(3C) is a more appropriate function because the transformation process occurs only once.</p> <p>The <code>strcoll()</code> function does not change the setting of <code>errno</code> if successful.</p> <p>Since no return value is reserved to indicate an error, an application wishing to check for error situations should set <code>errno</code> to 0, then call <code>strcoll()</code>, then check <code>errno</code>.</p>								
RETURN VALUES	<p>Upon successful completion, <code>strcoll()</code> returns an integer greater than, equal to, or less than zero in direct correlation to whether string <code>s1</code> is greater than, equal to, or less than the string <code>s2</code>. The comparison is based on strings interpreted as appropriate to the program's locale for category <code>LC_COLLATE</code> (see setlocale(3C)).</p> <p>On error, <code>strcoll()</code> may set <code>errno</code>, but no return value is reserved to indicate an error.</p>								
ERRORS	<p>The <code>strcoll()</code> function may fail if:</p> <table><tr><td><code>EINVAL</code></td><td>The <code>s1</code> or <code>s2</code> arguments contain characters outside the domain of the collating sequence.</td></tr></table>	<code>EINVAL</code>	The <code>s1</code> or <code>s2</code> arguments contain characters outside the domain of the collating sequence.						
<code>EINVAL</code>	The <code>s1</code> or <code>s2</code> arguments contain characters outside the domain of the collating sequence.								
FILES	<code>/usr/lib/locale/locale/locale.so.*</code> <code>LC_COLLATE</code> database for <i>locale</i>								
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>CSI</td><td>Enabled</td></tr><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe with exceptions</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	Standard	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
CSI	Enabled								
Interface Stability	Standard								
MT-Level	MT-Safe with exceptions								
SEE ALSO	localedef(1) , setlocale(3C) , string(3C) , strxfrm(3C) , wsxfrm(3C) , attributes(5) , environ(5) , standards(5)								

NAME	strerror, strerror_r – get error message string						
SYNOPSIS	<pre>#include <string.h> char *strerror(int <i>errnum</i>); int strerror_r(int <i>errnum</i>, char *<i>strerrbuf</i>, size_t <i>buflen</i>);</pre>						
DESCRIPTION	<p>The <code>strerror()</code> function maps the error number in <code>errnum</code> to an error message string, and returns a pointer to that string. It uses the same set of error messages as perror(3C). The returned string should not be overwritten.</p> <p>The <code>strerror_r()</code> function maps the error number in <code>errnum</code> to an error message string and returns the string in the buffer pointed to by <code>strerrbuf</code> with length <code>buflen</code>.</p>						
RETURN VALUES	<p>Upon successful completion, <code>strerror()</code> returns a pointer to the generated message string. Otherwise, it sets <code>errno</code> and returns a pointer to an error message string. It returns the string "Unknown error" if <code>errnum</code> is not a valid error number.</p> <p>Upon successful completion, <code>strerror_r()</code> returns 0. Otherwise it sets <code>errno</code> and returns the value of <code>errno</code> to indicate the error. It returns the string "Unknown error" in the buffer pointed to by <code>strerrbuf</code> if <code>errnum</code> is not a valid error number.</p>						
ERRORS	<p>These functions may fail if:</p> <p><code>EINVAL</code> The value of <code>errnum</code> is not a valid error number.</p> <p>The <code>strerror_r()</code> function may fail if:</p> <p><code>ERANGE</code> The <code>buflen</code> argument specifies insufficient storage to contain the generated message string.</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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	gettext(3C) , perror(3C) , setlocale(3C) , attributes(5) , standards(5)						
NOTES	If the application is linked with <code>-lintl</code> , then messages returned from these functions are in the native language specified by the <code>LC_MESSAGES</code> locale category. See setlocale(3C) .						

strfmon(3C)

NAME	strfmon – convert monetary value to string
SYNOPSIS	<pre>#include <monetary.h> ssize_t strfmon(char *restrict <i>s</i>, size_t <i>maxsize</i>, const char *restrict <i>format</i>, ...);</pre>
DESCRIPTION	<p>The <code>strfmon()</code> function places characters into the array pointed to by <i>s</i> as controlled by the string pointed to by <i>format</i>. No more than <i>maxsize</i> bytes are placed into the array.</p> <p>The format is a character string that contains two types of objects: plain characters, which are simply copied to the output stream, and conversion specifications, each of which results in the fetching of zero or more arguments which are converted and formatted. The results are undefined if there are insufficient arguments for the format. If the format is exhausted while arguments remain, the excess arguments are simply ignored.</p> <p>A conversion specification consists of the following sequence:</p> <ul style="list-style-type: none"> ■ a % character ■ optional flags ■ optional field width ■ optional left precision ■ optional right precision ■ a required conversion character that determines the conversion to be performed.
Flags	<p>One or more of the following optional flags can be specified to control the conversion:</p> <p><code>=<i>f</i></code> An <code>=</code> followed by a single character <i>f</i> which is used as the numeric fill character. The fill character must be representable in a single byte in order to work with precision and width counts. The default numeric fill character is the space character. This flag does not affect field width filling which always uses the space character. This flag is ignored unless a left precision (see below) is specified.</p> <p><code>^</code> Do not format the currency amount with grouping characters. The default is to insert the grouping characters if defined for the current locale.</p> <p><code>+ or (</code> Specify the style of representing positive and negative currency amounts. Only one of <code>'+'</code> or <code>'('</code> may be specified. If <code>'+'</code> is specified, the locale's equivalent of <code>+</code> and <code>'-'</code> are used (for example, in the U.S.A.: the empty string if positive and <code>'-'</code> if negative). If <code>'('</code> is specified, negative amounts are enclosed within parentheses. If neither flag is specified, the <code>'+'</code> style is used.</p> <p><code>!</code> Suppress the currency symbol from the output conversion.</p> <p><code>-</code> Specify the alignment. If this flag is present all fields are left-justified (padded to the right) rather than right-justified.</p>
Field Width	<p><code><i>w</i></code> A decimal digit string <i>w</i> specifying a minimum field width in bytes in which the result of the conversion is right-justified (or left-justified if the flag <code>'-'</code> is specified). The default is zero.</p>

Left Precision	<i>#n</i>	<p>A '#' followed by a decimal digit string <i>n</i> specifying a maximum number of digits expected to be formatted to the left of the radix character. This option can be used to keep the formatted output from multiple calls to the <code>strfmon()</code> aligned in the same columns. It can also be used to fill unused positions with a special character as in <code>\$***123.45</code>. This option causes an amount to be formatted as if it has the number of digits specified by <i>n</i>. If more than <i>n</i> digit positions are required, this conversion specification is ignored. Digit positions in excess of those actually required are filled with the numeric fill character (see the <code>=f</code> flag above).</p> <p>If grouping has not been suppressed with the '^' flag, and it is defined for the current locale, grouping separators are inserted before the fill characters (if any) are added. Grouping separators are not applied to fill characters even if the fill character is a digit.</p> <p>To ensure alignment, any characters appearing before or after the number in the formatted output such as currency or sign symbols are padded as necessary with space characters to make their positive and negative formats an equal length.</p>
Right Precision	<i>.p</i>	<p>A period followed by a decimal digit string <i>p</i> specifying the number of digits after the radix character. If the value of the right precision <i>p</i> is zero, no radix character appears. If a right precision is not included, a default specified by the current locale is used. The amount being formatted is rounded to the specified number of digits prior to formatting.</p>
Conversion Characters		<p>The conversion characters and their meanings are:</p> <p><i>i</i> The <code>double</code> argument is formatted according to the locale's international currency format (for example, in the U.S.A.: USD 1,234.56).</p> <p><i>n</i> The <code>double</code> argument is formatted according to the locale's national currency format (for example, in the U.S.A.: \$1,234.56).</p> <p><i>%</i> Convert to a %; no argument is converted. The entire conversion specification must be <code>%%</code>.</p>
Locale Information		<p>The <code>LC_MONETARY</code> category of the program's locale affects the behavior of this function including the monetary radix character (which may be different from the numeric radix character affected by the <code>LC_NUMERIC</code> category), the grouping separator, the currency symbols and formats. The international currency symbol should be in conformance with the ISO 4217: 1987 standard.</p>
RETURN VALUES		<p>If the total number of resulting bytes (including the terminating null byte) is not more than <i>maxsize</i>, <code>strfmon()</code> returns the number of bytes placed into the array pointed to by <i>s</i>, not including the terminating null byte. Otherwise, <code>-1</code> is returned, the contents of the array are indeterminate, and <code>errno</code> is set to indicate the error.</p>
ERRORS		<p>The <code>strfmon()</code> function will fail if:</p> <p><code>ENOSYS</code> The function is not supported.</p>

strfmon(3C)

E2BIG Conversion stopped due to lack of space in the buffer.

USAGE The behavior of `strfmon()` in an SUSv3-conforming application differs from its behavior in a non-conforming application as follows:

- With the conversion 'i', `strfmon()` uses information set to `int_p_cs_precedes`, `int_n_cs_precedes`, `int_p_sep_by_space`, `int_n_sep_by_space`, `int_p_sign_posn`, and `int_n_sign_posn` of the current locale instead of `p_cs_precedes`, `n_cs_precedes`, `p_sep_by_space`, `n_sep_by_space`, `p_sign_posn`, and `n_sign_posn`, respectively.
- With the conversion 'i', `strfmon()` uses the fourth character of the string set to `int_curr_symbol` of the current locale instead of a space for `int_p_sep_by_space` and `int_n_sep_by_space`.
- When the value of `p_sep_by_space`, `n_sep_by_space`, `int_p_sep_by_space`, or `int_n_sep_by_space` is set to 2 in the current locale, `strfmon()` separates the currency symbol from the sign string by a space, if adjacent; otherwise, `strfmon()` separates the sign string from the value by a space.

EXAMPLES **EXAMPLE 1** A sample output of `strfmon()`.

Given a locale for the U.S.A. and the values 123.45, -123.45, and 3456.781:

Conversion Specification	Output	Comments
%n	\$123.45 -\$123.45 \$3,456.78	default formatting
%11n	\$123.45 -\$123.45 \$3,456.78	right align within an 11 character field
%#5n	\$123.45 -\$123.45 \$3,456.78	aligned columns for values up to 99,999
%=#5n	\$\$\$123.45 -\$\$\$\$123.45 \$*3,456.78	specify a fill character
%=0#5n	\$000123.45	fill characters do not use

EXAMPLE 1 A sample output of `strfmon()`. (Continued)

Conversion Specification	Output	Comments
	-\$000123.45	grouping even if the fill character is a digit
	\$03,456.78	
%^#5n	\$123.45	disable the grouping separator
	-\$123.45	
	\$3456.78	
%^#5.0n	\$123	round off to whole units
	-\$123	
	\$3457	
%^#5.4n	\$123.4500	increase the precision
	-\$123.4500	
	\$3456.7810	
%(#5n	123.45	use an alternative pos/neg style
	(\$123.45)	
	\$3,456.78	
%!(#5n	123.45	disable the currency symbol
	(123.45)	
	3,456.78	

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

The `strfmon()` function can be used safely in multithreaded applications, as long as `setlocale(3C)` is not called to change the locale.

SEE ALSO `localeconv(3C)`, `setlocale(3C)`, `attributes(5)`, `standards(5)`

strftime(3C)

NAME	strftime, cftime, ascftime – convert date and time to string																						
SYNOPSIS	<pre>#include <time.h> size_t strftime(char *restrict s, size_t <i>maxsize</i>, const char *restrict <i>format</i>, const struct tm *restrict <i>timeptr</i>); int cftime(char *s, char *<i>format</i>, const time_t *<i>clock</i>); int ascftime(char *s, const char *<i>format</i>, const struct tm *<i>timeptr</i>);</pre>																						
DESCRIPTION	<p>The <code>strftime()</code>, <code>ascftime()</code>, and <code>cftime()</code> functions place bytes into the array pointed to by <code>s</code> as controlled by the string pointed to by <code>format</code>. The <code>format</code> string consists of zero or more conversion specifications and ordinary characters. A conversion specification consists of a '%' (percent) character and one or two terminating conversion characters that determine the conversion specification's behavior. All ordinary characters (including the terminating null byte) are copied unchanged into the array pointed to by <code>s</code>. If copying takes place between objects that overlap, the behavior is undefined. For <code>strftime()</code>, no more than <code>maxsize</code> bytes are placed into the array.</p> <p>If <code>format</code> is <code>(char *)0</code>, then the locale's default format is used. For <code>strftime()</code> the default format is the same as <code>%c</code>; for <code>cftime()</code> and <code>ascftime()</code> the default format is the same as <code>%C</code>. <code>cftime()</code> and <code>ascftime()</code> first try to use the value of the environment variable <code>CFTIME</code>, and if that is undefined or empty, the default format is used.</p> <p>Each conversion specification is replaced by appropriate characters as described in the following list. The appropriate characters are determined by the <code>LC_TIME</code> category of the program's locale and by the values contained in the structure pointed to by <code>timeptr</code> for <code>strftime()</code> and <code>ascftime()</code>, and by the time represented by <code>clock</code> for <code>cftime()</code>.</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 10px;">%%</td> <td>Same as %.</td> </tr> <tr> <td style="padding-right: 10px;">%a</td> <td>Locale's abbreviated weekday name.</td> </tr> <tr> <td style="padding-right: 10px;">%A</td> <td>Locale's full weekday name.</td> </tr> <tr> <td style="padding-right: 10px;">%b</td> <td>Locale's abbreviated month name.</td> </tr> <tr> <td style="padding-right: 10px;">%B</td> <td>Locale's full month name.</td> </tr> <tr> <td style="padding-right: 10px;">%c</td> <td>Locale's appropriate date and time representation.</td> </tr> <tr> <td style="padding-right: 10px;">Default %C</td> <td>Locale's date and time representation as produced by <code>date(1)</code>.</td> </tr> <tr> <td style="padding-right: 10px;">Standard conforming %C</td> <td>Century number (the year divided by 100 and truncated to an integer as a decimal number [1,99]); single digits are preceded by 0; see <code>standards(5)</code>.</td> </tr> <tr> <td style="padding-right: 10px;">%d</td> <td>Day of month [1,31]; single digits are preceded by 0.</td> </tr> <tr> <td style="padding-right: 10px;">%D</td> <td>Date as <code>%m/%d/%y</code>.</td> </tr> <tr> <td style="padding-right: 10px;">%e</td> <td>Day of month [1,31]; single digits are preceded by a space.</td> </tr> </table>	%%	Same as %.	%a	Locale's abbreviated weekday name.	%A	Locale's full weekday name.	%b	Locale's abbreviated month name.	%B	Locale's full month name.	%c	Locale's appropriate date and time representation.	Default %C	Locale's date and time representation as produced by <code>date(1)</code> .	Standard conforming %C	Century number (the year divided by 100 and truncated to an integer as a decimal number [1,99]); single digits are preceded by 0; see <code>standards(5)</code> .	%d	Day of month [1,31]; single digits are preceded by 0.	%D	Date as <code>%m/%d/%y</code> .	%e	Day of month [1,31]; single digits are preceded by a space.
%%	Same as %.																						
%a	Locale's abbreviated weekday name.																						
%A	Locale's full weekday name.																						
%b	Locale's abbreviated month name.																						
%B	Locale's full month name.																						
%c	Locale's appropriate date and time representation.																						
Default %C	Locale's date and time representation as produced by <code>date(1)</code> .																						
Standard conforming %C	Century number (the year divided by 100 and truncated to an integer as a decimal number [1,99]); single digits are preceded by 0; see <code>standards(5)</code> .																						
%d	Day of month [1,31]; single digits are preceded by 0.																						
%D	Date as <code>%m/%d/%y</code> .																						
%e	Day of month [1,31]; single digits are preceded by a space.																						

%F	Equivalent to %Y-%m-%d (the ISO 8601:2000 standard date format).
%g	Week-based year within century [00,99].
%G	Week-based year, including the century [0000,9999].
%h	Locale's abbreviated month name.
%H	Hour (24-hour clock) [0,23]; single digits are preceded by 0.
%I	Hour (12-hour clock) [1,12]; single digits are preceded by 0.
%j	Day number of year [1,366]; single digits are preceded by 0.
%k	Hour (24-hour clock) [0,23]; single digits are preceded by a blank.
%l	Hour (12-hour clock) [1,12]; single digits are preceded by a blank.
%m	Month number [1,12]; single digits are preceded by 0.
%M	Minute [00,59]; leading 0 is permitted but not required.
%n	Insert a NEWLINE.
%p	Locale's equivalent of either a.m. or p.m.
%r	Appropriate time representation in 12-hour clock format with %p.
%R	Time as %H:%M.
%S	Seconds [00,60]; the range of values is [00,60] rather than [00,59] to allow for the occasional leap second and even more occasional double leap second.
%t	Insert a TAB.
%T	Time as %H:%M:%S.
%u	Weekday as a decimal number [1,7], with 1 representing Monday. See NOTES below.
%U	Week number of year as a decimal number [00,53], with Sunday as the first day of week 1.
%V	The ISO 8601 week number as a decimal number [01,53]. In the ISO 8601 week-based system, weeks begin on a Monday and week 1 of the year is the week that includes both January 4th and the first Thursday of the year. If the first Monday of January is the 2nd, 3rd, or 4th, the preceding days are part of the last week of the preceding year. See NOTES below.
%w	Weekday as a decimal number [0,6], with 0 representing Sunday.
%W	Week number of year as a decimal number [00,53], with Monday as the first day of week 1.
%x	Locale's appropriate date representation.
%X	Locale's appropriate time representation.

strptime(3C)

%Y	Year within century [00,99].
%Y	Year, including the century (for example 1993).
%z	Replaced by offset from UTC in ISO 8601:2000 standard format (+hhmm or -hhmm), or by no characters if no timezone is determinable. For example, "-0430" means 4 hours 30 minutes behind UTC (west of Greenwich). If <code>tm_isdst</code> is zero, the standard time offset is used. If <code>tm_isdst</code> is greater than zero, the daylight savings time offset is used. If <code>tm_isdst</code> is negative, no characters are returned.
%Z	Time zone name or abbreviation, or no bytes if no time zone information exists.

If a conversion specification does not correspond to any of the above or to any of the modified conversion specifications listed below, the behavior is undefined and 0 is returned.

The difference between %U and %W (and also between modified conversion specifications %OU and %OW) lies in which day is counted as the first of the week. Week number 1 is the first week in January starting with a Sunday for %U or a Monday for %W. Week number 0 contains those days before the first Sunday or Monday in January for %U and %W, respectively.

Modified Conversion Specifications

Some conversion specifications can be modified by the E and O modifiers to indicate that an alternate format or specification should be used rather than the one normally used by the unmodified conversion specification. If the alternate format or specification does not exist in the current locale, the behavior will be as if the unmodified specification were used.

%Ec	Locale's alternate appropriate date and time representation.
%EC	Name of the base year (period) in the locale's alternate representation.
%Eg	Offset from %EC of the week-based year in the locale's alternative representation.
%EG	Full alternative representation of the week-based year.
%Ex	Locale's alternate date representation.
%EX	Locale's alternate time representation.
%Ey	Offset from %EC (year only) in the locale's alternate representation.
%EY	Full alternate year representation.
%Od	Day of the month using the locale's alternate numeric symbols.
%Oe	Same as %Od.
%Og	Week-based year (offset from %C) in the locale's alternate representation and using the locale's alternate numeric symbols.
%OH	Hour (24-hour clock) using the locale's alternate numeric symbols.

%OI	Hour (12-hour clock) using the locale's alternate numeric symbols.
%Om	Month using the locale's alternate numeric symbols.
%OM	Minutes using the locale's alternate numeric symbols.
%OS	Seconds using the locale's alternate numeric symbols.
%Ou	Weekday as a number in the locale's alternate numeric symbols.
%OU	Week number of the year (Sunday as the first day of the week) using the locale's alternate numeric symbols.
%Ow	Number of the weekday (Sunday=0) using the locale's alternate numeric symbols.
%OW	Week number of the year (Monday as the first day of the week) using the locale's alternate numeric symbols.
%Oy	Year (offset from %C) in the locale's alternate representation and using the locale's alternate numeric symbols.

Selecting the Output Language

By default, the output of `strptime()`, `cftime()`, and `ascftime()` appear in U.S. English. The user can request that the output of `strptime()`, `cftime()`, or `ascftime()` be in a specific language by setting the `LC_TIME` category using `setlocale()`.

Time Zone

Local time zone information is used as though `tzset(3C)` were called.

RETURN VALUES

The `strptime()`, `cftime()`, and `ascftime()` functions return the number of characters placed into the array pointed to by `s`, not including the terminating null character. If the total number of resulting characters including the terminating null character is more than `maxsize`, `strptime()` returns 0 and the contents of the array are indeterminate.

EXAMPLES

EXAMPLE 1 An example of the `strptime()` function.

The following example illustrates the use of `strptime()` for the POSIX locale. It shows what the string in `str` would look like if the structure pointed to by `tmpr` contains the values corresponding to Thursday, August 28, 1986 at 12:44:36.

```
strptime (str, strsize, "%A %b %d %j", tmpr)
```

This results in `str` containing "Thursday Aug 28 240".

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe
CSI	Enabled

strptime(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	strptime() is Standard.

SEE ALSO [date\(1\)](#), [ctime\(3C\)](#), [mktime\(3C\)](#), [setlocale\(3C\)](#), [strptime\(3C\)](#), [tzset\(3C\)](#), [TIMEZONE\(4\)](#), [zoneinfo\(4\)](#), [attributes\(5\)](#), [environ\(5\)](#), [standards\(5\)](#)

NOTES The conversion specification for %V was changed in the Solaris 7 release. This change was based on the public review draft of the ISO C9x standard at that time. Previously, the specification stated that if the week containing 1 January had fewer than four days in the new year, it became week 53 of the previous year. The ISO C9x standard committee subsequently recognized that that specification had been incorrect.

The conversion specifications for %g, %G, %Eg, %EG, and %Og were added in the Solaris 7 release. This change was based on the public review draft of the ISO C9x standard at that time. These specifications are evolving. If the ISO C9x standard is finalized with a different conclusion, these specifications will change to conform to the ISO C9x standard decision.

The conversion specification for %u was changed in the Solaris 8 release. This change was based on the XPG4 specification.

If using the %Z specifier and zoneinfo timezones and if the input date is outside the range 20:45:52 UTC, December 13, 1901 to 03:14:07 UTC, January 19, 2038, the timezone name may not be correct.

NAME	string, strcasecmp, strncasecmp, strcat, strncat, strlcat, strchr, strrchr, strcmp, strncmp, strcpy, strncpy, strlcpy, strcspn, strspn, strdup, strlen, strpbrk, strstr, strtok, strtok_r – string operations
SYNOPSIS	<pre> #include <strings.h> int strcasecmp(const char *s1, const char *s2); int strncasecmp(const char *s1, const char *s2, size_t n); #include <string.h> char *strcat(char *restrict s1, const char *restrict s2); char *strncat(char *restrict s1, const char *restrict s2, size_t n); size_t strlcat(char *dst, const char *src, size_t dstsize); char *strchr(const char *s, int c); char *strrchr(const char *s, int c); int strcmp(const char *s1, const char *s2); int strncmp(const char *s1, const char *s2, size_t n); char *strcpy(char *restrict s1, const char *restrict s2); char *strncpy(char *restrict s1, const char *restrict s2, size_t n); size_t strlcpy(char *dst, const char *src, size_t dstsize); size_t strcspn(const char *s1, const char *s2); size_t strspn(const char *s1, const char *s2); char *strdup(const char *s1); size_t strlen(const char *s); char *strpbrk(const char *s1, const char *s2); char *strstr(const char *s1, const char *s2); char *strtok(char *restrict s1, const char *restrict s2); char *strtok_r(char *s1, const char *s2, char **lasts); ISO C++ #include <string.h> const char *strchr(const char *s, int c); const char *strpbrk(const char *s1, const char *s2); const char *strrchr(const char *s, int c); const char *strstr(const char *s1, const char *s2); #include <cstring> </pre>

string(3C)

	<pre>char *std::strchr(char *s, int c); char *std::strpbrk(char *s1, const char *s2); char *std::strrchr(char *s, int c); char *std::strstr(char *s1, const char *s2);</pre>
DESCRIPTION	<p>The arguments <i>s</i>, <i>s1</i>, and <i>s2</i> point to strings (arrays of characters terminated by a null character). The <code>strcat()</code>, <code>strncat()</code>, <code>strlcat()</code>, <code>strcpy()</code>, <code>strncpy()</code>, <code>strncpy()</code>, <code>strtok()</code>, and <code>strtok_r()</code> functions all alter their first argument. These functions do not check for overflow of the array pointed to by the first argument.</p>
strcasecmp() , strncasecmp()	<p>The <code>strcasecmp()</code> and <code>strncasecmp()</code> functions are case-insensitive versions of <code>strcmp()</code> and <code>strncmp()</code> respectively, described below. They assume the ASCII character set and ignore differences in case when comparing lower and upper case characters.</p>
strcat() , strncat() , strlcat()	<p>The <code>strcat()</code> function appends a copy of string <i>s2</i>, including the terminating null character, to the end of string <i>s1</i>. The <code>strncat()</code> function appends at most <i>n</i> characters. Each returns a pointer to the null-terminated result. The initial character of <i>s2</i> overrides the null character at the end of <i>s1</i>.</p> <p>The <code>strlcat()</code> function appends at most $(dstsize - strlen(dst) - 1)$ characters of <i>src</i> to <i>dst</i> (<i>dstsize</i> being the size of the string buffer <i>dst</i>). If the string pointed to by <i>dst</i> contains a null-terminated string that fits into <i>dstsize</i> bytes when <code>strlcat()</code> is called, the string pointed to by <i>dst</i> will be a null-terminated string that fits in <i>dstsize</i> bytes (including the terminating null character) when it completes, and the initial character of <i>src</i> will override the null character at the end of <i>dst</i>. If the string pointed to by <i>dst</i> is longer than <i>dstsize</i> bytes when <code>strlcat()</code> is called, the string pointed to by <i>dst</i> will not be changed. The function returns $\min\{dstsize, strlen(dst)\} + strlen(src)$. Buffer overflow can be checked as follows:</p> <pre>if (strlcat(dst, src, dstsize) >= dstsize) return -1;</pre>
strchr() , strrchr()	<p>The <code>strchr()</code> function returns a pointer to the first occurrence of <i>c</i> (converted to a char) in string <i>s</i>, or a null pointer if <i>c</i> does not occur in the string. The <code>strrchr()</code> function returns a pointer to the last occurrence of <i>c</i>. The null character terminating a string is considered to be part of the string.</p>
strcmp() , strncmp()	<p>The <code>strcmp()</code> function compares two strings byte-by-byte, according to the ordering of your machine's character set. The function returns an integer greater than, equal to, or less than 0, if the string pointed to by <i>s1</i> is greater than, equal to, or less than the string pointed to by <i>s2</i> respectively. The sign of a non-zero return value is determined by the sign of the difference between the values of the first pair of bytes that differ in the strings being compared. The <code>strncmp()</code> function makes the same comparison but looks at a maximum of <i>n</i> bytes. Bytes following a null byte are not compared.</p>

strcpy() , strncpy() , strncpy() , strncpy()	<p>The <code>strcpy()</code> function copies string <code>s2</code> to <code>s1</code>, including the terminating null character, stopping after the null character has been copied. The <code>strncpy()</code> function copies exactly <code>n</code> bytes, truncating <code>s2</code> or adding null characters to <code>s1</code> if necessary. The result will not be null-terminated if the length of <code>s2</code> is <code>n</code> or more. Each function returns <code>s1</code>.</p> <p>The <code>strncpy()</code> function copies at most <code>dstsize-1</code> characters (<code>dstsize</code> being the size of the string buffer <code>dst</code>) from <code>src</code> to <code>dst</code>, truncating <code>src</code> if necessary. The result is always null-terminated. The function returns <code>strlen(src)</code>. Buffer overflow can be checked as follows:</p> <pre>if (strncpy(dst, src, dstsize) >= dstsize) return -1;</pre>
strcspn() , strcspn()	<p>The <code>strcspn()</code> function returns the length of the initial segment of string <code>s1</code> that consists entirely of characters not from string <code>s2</code>. The <code>strspn()</code> function returns the length of the initial segment of string <code>s1</code> that consists entirely of characters from string <code>s2</code>.</p>
strdup()	<p>The <code>strdup()</code> function returns a pointer to a new string that is a duplicate of the string pointed to by <code>s1</code>. The returned pointer can be passed to <code>free()</code>. The space for the new string is obtained using <code>malloc(3C)</code>. If the new string cannot be created, a null pointer is returned and <code>errno</code> may be set to <code>ENOMEM</code> to indicate that the storage space available is insufficient.</p>
strlen()	<p>The <code>strlen()</code> function returns the number of bytes in <code>s</code>, not including the terminating null character.</p>
strpbrk()	<p>The <code>strpbrk()</code> function returns a pointer to the first occurrence in string <code>s1</code> of any character from string <code>s2</code>, or a null pointer if no character from <code>s2</code> exists in <code>s1</code>.</p>
strstr()	<p>The <code>strstr()</code> function locates the first occurrence of the string <code>s2</code> (excluding the terminating null character) in string <code>s1</code> and returns a pointer to the located string, or a null pointer if the string is not found. If <code>s2</code> points to a string with zero length (that is, the string ""), the function returns <code>s1</code>.</p>
strtok()	<p>A sequence of calls to <code>strtok()</code> breaks the string pointed to by <code>s1</code> into a sequence of tokens, each of which is delimited by a byte from the string pointed to by <code>s2</code>. The first call in the sequence has <code>s1</code> as its first argument, and is followed by calls with a null pointer as their first argument. The separator string pointed to by <code>s2</code> can be different from call to call.</p> <p>The first call in the sequence searches the string pointed to by <code>s1</code> for the first byte that is not contained in the current separator string pointed to by <code>s2</code>. If no such byte is found, then there are no tokens in the string pointed to by <code>s1</code> and <code>strtok()</code> returns a null pointer. If such a byte is found, it is the start of the first token.</p> <p>The <code>strtok()</code> function then searches from there for a byte that is contained in the current separator string. If no such byte is found, the current token extends to the end of the string pointed to by <code>s1</code>, and subsequent searches for a token return a null pointer. If such a byte is found, it is overwritten by a null byte that terminates the current token. The <code>strtok()</code> function saves a pointer to the following byte in thread-specific data, from which the next search for a token starts.</p>

string(3C)

Each subsequent call, with a null pointer as the value of the first argument, starts searching from the saved pointer and behaves as described above.

See Example 1, 2, and 3 in the EXAMPLES section for examples of `strtok()` usage.

`strtok_r()`

The `strtok_r()` function considers the null-terminated string *s1* as a sequence of zero or more text tokens separated by spans of one or more characters from the separator string *s2*. The argument *lasts* points to a user-provided pointer which points to stored information necessary for `strtok_r()` to continue scanning the same string.

In the first call to `strtok_r()`, *s1* points to a null-terminated string, *s2* to a null-terminated string of separator characters, and the value pointed to by *lasts* is ignored. The `strtok_r()` function returns a pointer to the first character of the first token, writes a null character into *s1* immediately following the returned token, and updates the pointer to which *lasts* points.

In subsequent calls, *s1* is a null pointer and *lasts* is unchanged from the previous call so that subsequent calls move through the string *s1*, returning successive tokens until no tokens remain. The separator string *s2* can be different from call to call. When no token remains in *s*, a null pointer is returned.

See Example 3 in the EXAMPLES section for an example of `strtok_r()` usage.

EXAMPLES

EXAMPLE 1 Search for word separators.

The following example searches for tokens separated by space characters.

```
#include <string.h>
...
char *token;
char *line = "LINE TO BE SEPARATED";
char *search = " ";

/* Token will point to "LINE". */
token = strtok(line, search);

/* Token will point to "TO". */
token = strtok(NULL, search);
```

EXAMPLE 2 Break a Line.

The following example uses `strtok` to break a line into two character strings separated by any combination of SPACES, TABS, or NEWLINES.

```
#include <string.h>
...
struct element {
    char *key;
    char *data;
};
...
char line[LINE_MAX];
```

EXAMPLE 2 Break a Line. (Continued)

```
char *key, *data;
...
key = strtok(line, " \n");
data = strtok(NULL, " \n");
```

EXAMPLE 3 Search for tokens.

The following example uses both `strtok()` and `strtok_r()` to search for tokens separated by one or more characters from the string pointed to by the second argument, `"/"`.

```
#define __EXTENSIONS__
#include <stdio.h>
#include <string.h>

int
main() {
    char *buf="5/90/45";
    char *token;
    char *lasts;

    printf("tokenizing \"%s\" with strtok():\n", buf);
    if ((token = strtok(buf, "/")) != NULL) {
        printf("token = \"%s\"\n", token);
        while ((token = strtok(NULL, "/")) != NULL) {
            printf("token = \"%s\"\n", token);
        }
    }

    buf = "//5//90//45//";
    printf("\ntokenizing \"%s\" with strtok_r():\n", buf);
    if ((token = strtok_r(buf, "/", &lasts)) != NULL) {
        printf("token = \"%s\"\n", token);
        while ((token = strtok_r(NULL, "/", &lasts)) != NULL) {
            printf("token = \"%s\"\n", token);
        }
    }
}
```

When compiled and run, this example produces the following output:

```
tokenizing "5/90/45" with strtok():
token = "5"
token = "90"
token = "45"

tokenizing "//5//90//45//" with strtok_r():
token = "5"
token = "90"
token = "45"
```

string(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	See below.
MT-Level	See below.

The `strlcat()` and `strncpy()` functions are Stable. The remaining functions are Standard.

The `strtok()` and `strdup()` functions are MT-Safe. The remaining functions are Async-Signal-Safe.

SEE ALSO `malloc(3C)`, `setlocale(3C)`, `strxfrm(3C)`, `attributes(5)`, `standards(5)`

NOTES When compiling multithreaded applications, the `_REENTRANT` flag must be defined on the compile line. This flag should only be used in multithreaded applications.

A single-threaded application can gain access to `strtok_r()` only by defining `__EXTENSIONS__` or by defining `_POSIX_C_SOURCE` to a value greater than or equal to 199506L.

All of these functions assume the default locale "C." For some locales, `strxfrm(3C)` should be applied to the strings before they are passed to the functions.

The `strtok()` function is safe to use in multithreaded applications because it saves its internal state in a thread-specific data area. However, its use is discouraged, even for single-threaded applications. The `strtok_r()` function should be used instead.

NAME	string_to_decimal, file_to_decimal, func_to_decimal – parse characters into decimal record								
SYNOPSIS	<pre>#include <floatingpoint.h> void string_to_decimal(char **pc, int nmax, int fortran_conventions, decimal_record *pd, enum decimal_string_form *pform, char **pechar); void func_to_decimal(char **pc, int nmax, int fortran_conventions, decimal_record *pd, enum decimal_string_form *pform, char **pechar, int (*pget)(void), int *pnread, int (*punget)(int c)); #include <stdio.h> void file_to_decimal(char **pc, int nmax, int fortran_conventions, decimal_record *pd, enum decimal_string_form *pform, char **pechar, FILE *pf, int *pnread);</pre>								
DESCRIPTION	<p>These functions attempt to parse a numeric token from at most <i>nmax</i> characters read from a string <i>**pc</i>, a file <i>*pf</i>, or function (<i>*pget</i>) (). They set the decimal record <i>*pd</i> to reflect the value of the numeric token recognized and set <i>*pform</i> and <i>*pechar</i> to indicate its form.</p> <p>The accepted forms for the numeric token consist of an initial, possibly empty, sequence of white-space characters, as defined by isspace(3C), followed by a subject sequence representing a numeric value, infinity, or NaN. The subject sequence consists of an optional plus or minus sign followed by one of the following:</p> <ul style="list-style-type: none"> ■ a non-empty sequence of decimal digits optionally containing a decimal point character, then an optional exponent part ■ one of INF or INFINITY, ignoring case ■ one of NAN or NAN(<i>string</i>), ignoring case in the NAN part; <i>string</i> can be any sequence of characters not containing ')' (right parenthesis) or '\0' (null). <p>The <i>fortran_conventions</i> argument provides additional control over the set of accepted forms. It must be one of the following values:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">0</td> <td>no Fortran conventions</td> </tr> <tr> <td>1</td> <td>Fortran list-directed input conventions</td> </tr> <tr> <td>2</td> <td>Fortran formatted input conventions, blanks are ignored</td> </tr> <tr> <td>3</td> <td>Fortran formatted input conventions, blanks are interpreted as zeroes</td> </tr> </table> <p>When <i>fortran_conventions</i> is zero, the decimal point character is the current locale's decimal point character, and the exponent part consists of the letter E or e followed by an optional sign and a non-empty string of decimal digits.</p> <p>When <i>fortran_conventions</i> is non-zero, the decimal point character is "." (period), and the exponent part consists of either a sign or one of the letters E, e, D, d, Q, or q followed by an optional sign, then a non-empty string of decimal digits.</p>	0	no Fortran conventions	1	Fortran list-directed input conventions	2	Fortran formatted input conventions, blanks are ignored	3	Fortran formatted input conventions, blanks are interpreted as zeroes
0	no Fortran conventions								
1	Fortran list-directed input conventions								
2	Fortran formatted input conventions, blanks are ignored								
3	Fortran formatted input conventions, blanks are interpreted as zeroes								

string_to_decimal(3C)

When *fortran_conventions* is 2 or 3, blanks can appear in the digit strings for the integer, fraction, and exponent parts, between the exponent delimiter and optional exponent sign, and after an INF, INFINITY, NAN, or NAN(*string*). When *fortran_conventions* is 2, all blanks are ignored. When *fortran_conventions* is 3, blanks in digit strings are interpreted as zeros and other blanks are ignored.

The following table summarizes the accepted forms and shows the corresponding values to which **pform* and *pd->fpclass* are set. Here *digits* represents any string of decimal digits, "." (period) stands for the decimal point character, and *exponent* represents the exponent part as defined above. Numbers in brackets refer to the notes following the table.

form	*pform	pd->fpclass
all white space [1]	whitespace_form	fp_zero
<i>digits</i>	fixed_int_form	fp_normal [2]
<i>digits.</i>	fixed_intdot_form	fp_normal [2]
<i>.digits</i>	fixed_dotfrac_form	fp_normal [2]
<i>digits.digits</i>	fixed_intdotfrac_form	fp_normal [2]
<i>digits exponent</i>	floating_int_form	fp_normal [2]
<i>digits. exponent</i>	floating_intdot_form	fp_normal [2]
<i>.digits exponent</i>	floating_dotfrac_form	fp_normal [2]
<i>digits.digits exponent</i>	floating_intdotfrac_form	fp_normal [2]
INF	inf_form	fp_infinity
INFINITY	infinity_form	fp_infinity
NAN	nan_form	fp_quiet
NAN(<i>string</i>)	nanstring_form	fp_quiet
none of the above	invalid_form	fp_signaling

Notes:

1. The *whitespace_form* is accepted only when *fortran_conventions* is 2 or 3 and is interpreted as zero.
2. For all numeric forms, *pd->fpclass* is set to *fp_normal* if any non-zero digits appear in the integer or fraction parts, and otherwise *pd->fpclass* is set to *fp_zero*.

If the accepted token has one of the numeric forms and represents a non-zero number x , its significant digits are stored in `pd->ds`. Leading and trailing zeroes and the radix point are omitted. `pd->sign` and `pd->exponent` are set so that if m is the integer represented by `pd->ds`,

$$-1^{**}(\text{pd->sign}) * m * 10^{**}(\text{pd->exponent})$$

approximates x to at least 511 significant digits. `pd->more` is set to 1 if this approximation is not exact (that is, the accepted token contains additional non-zero digits beyond those copied to `pd->ds`) and to 0 otherwise.

If the accepted token has the NAN(*string*) form, up to 511 characters from the string part are copied to `pd->ds`.

`pd->ds` is always terminated by a null byte, and `pd->ndigits` is set to the length of the string stored in `pd->ds`.

On entry, `*pc` points to the beginning of a character string buffer. The `string_to_decimal()` function reads characters from this buffer until either enough characters are read to delimit the accepted token (for example, a null character marking the end of the string is found) or the limit of `nmax` characters is reached. The `file_to_decimal()` function reads characters from the file `*pf` and stores them in the buffer. The `func_to_decimal()` function reads characters one at a time by calling the function (`*pget`)() and stores them in the buffer; (`*pget`)() must return integer values in the range -1 to 255, where -1 is interpreted as EOF and 0, ..., 255 are interpreted as unsigned char values. Both `file_to_decimal()` and `func_to_decimal()` read characters until either enough characters are read to delimit the accepted token, EOF is encountered, or the limit of `nmax` characters is reached. These functions, therefore, typically read one or more additional characters beyond the end of the accepted token and attempt to push back any excess characters read. Provided that the `punget` argument is not NULL, `func_to_decimal()` pushes back characters one at a time by calling (`*punget`)(*c*), where *c* is an integer in the range 0 to 255 corresponding to a value previously read via (`*pget`)(). After pushing back as many excess characters as possible, `file_to_decimal()` and `func_to_decimal()` store a null byte in the buffer following the last character read and not pushed back and set `*pnread` to the number of characters stored in the buffer prior to this null byte. Since these functions can read up to `nmax` characters, the buffer must be large enough to hold `nmax + 1`.

On exit, `*pc` points to the next character in the buffer past the last one that was accepted as part of the numeric token. If no valid token is found, `*pc` is unchanged. If `file_to_decimal()` and `func_to_decimal()` successfully push back all unused characters, `*pc` points to the null byte stored in the buffer following the last character read and not pushed back.

If the accepted token contains an exponent part, `*pechar` is set to point to the position in the buffer where the first character of the exponent field is stored. If the accepted token does not contain an exponent part, `*pechar` is set to NULL.

string_to_decimal(3C)

USAGE | If the `_IOWRT` flag is set in `*pf`, `file_to_decimal()` reads characters directly from the file buffer until a null character is found. (The `_IOWRT` flag should only be set when `file_to_decimal()` is called from `sscanf(3C)`.) Otherwise, `file_to_decimal()` uses `getc_unlocked(3C)`, so it is not MT-safe unless the caller holds the stream lock.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe with exceptions

SEE ALSO | `ctype(3C)`, `decimal_to_floating(3C)`, `getc_unlocked(3C)`, `isspace(3C)`, `localeconv(3C)`, `scanf(3C)`, `setlocale(3C)`, `strtod(3C)`, `ungetc(3C)`, `attributes(5)`

NAME	strptime – date and time conversion																		
SYNOPSIS	<pre>#include <time.h> char *strptime(const char *restrict <i>buf</i>, const char *restrict <i>format</i>, struct tm *restrict <i>tm</i>);</pre>																		
Non-zeroing Behavior	<pre>cc [<i>flag... </i>] <i>file... </i> -D_STRPTIME_DONTZERO [<i>library... </i>] #include <time.h> char *strptime(const char *restrict <i>buf</i>, const char *restrict <i>format</i>, struct tm *restrict <i>tm</i>);</pre>																		
DESCRIPTION	<p>The <code>strptime()</code> function converts the character string pointed to by <i>buf</i> to values which are stored in the <code>tm</code> structure pointed to by <i>tm</i>, using the format specified by <i>format</i>.</p> <p>The <i>format</i> argument is composed of zero or more conversion specifications. Each conversion specification is composed of a “%” (percent) character followed by one or two conversion characters which specify the replacement required. One or more white space characters (as specified by isspace(3C)) may precede or follow a conversion specification. There must be white-space or other non-alphanumeric characters between any two conversion specifications.</p> <p>A non-zeroing version of <code>strptime()</code>, described below under Non-zeroing Behavior, is provided if <code>_STRPTIME_DONTZERO</code> is defined.</p>																		
Conversion Specifications	<p>The following conversion specifications are supported:</p> <table border="0"> <tr> <td style="padding-right: 20px;">%%</td> <td>Same as %.</td> </tr> <tr> <td>%a</td> <td>Day of week, using the locale’s weekday names; either the abbreviated or full name may be specified.</td> </tr> <tr> <td>%A</td> <td>Same as %a.</td> </tr> <tr> <td>%b</td> <td>Month, using the locale’s month names; either the abbreviated or full name may be specified.</td> </tr> <tr> <td>%B</td> <td>Same as %b.</td> </tr> <tr> <td>%c</td> <td>Locale’s appropriate date and time representation.</td> </tr> <tr> <td>%C</td> <td>Century number (the year divided by 100 and truncated to an integer as a decimal number [1,99]); single digits are preceded by 0. If %C is used without the %y specifier, <code>strptime()</code> assumes the year offset is zero in whichever century is specified. Note the behavior of %C in the absence of %y is not specified by any of the standards or specifications described on the standards(5) manual page, so portable applications should not depend on it. This behavior may change in a future release.</td> </tr> <tr> <td>%d</td> <td>Day of month [1,31]; leading zero is permitted but not required.</td> </tr> <tr> <td>%D</td> <td>Date as %m/%d/%y.</td> </tr> </table>	%%	Same as %.	%a	Day of week, using the locale’s weekday names; either the abbreviated or full name may be specified.	%A	Same as %a.	%b	Month, using the locale’s month names; either the abbreviated or full name may be specified.	%B	Same as %b.	%c	Locale’s appropriate date and time representation.	%C	Century number (the year divided by 100 and truncated to an integer as a decimal number [1,99]); single digits are preceded by 0. If %C is used without the %y specifier, <code>strptime()</code> assumes the year offset is zero in whichever century is specified. Note the behavior of %C in the absence of %y is not specified by any of the standards or specifications described on the standards(5) manual page, so portable applications should not depend on it. This behavior may change in a future release.	%d	Day of month [1,31]; leading zero is permitted but not required.	%D	Date as %m/%d/%y.
%%	Same as %.																		
%a	Day of week, using the locale’s weekday names; either the abbreviated or full name may be specified.																		
%A	Same as %a.																		
%b	Month, using the locale’s month names; either the abbreviated or full name may be specified.																		
%B	Same as %b.																		
%c	Locale’s appropriate date and time representation.																		
%C	Century number (the year divided by 100 and truncated to an integer as a decimal number [1,99]); single digits are preceded by 0. If %C is used without the %y specifier, <code>strptime()</code> assumes the year offset is zero in whichever century is specified. Note the behavior of %C in the absence of %y is not specified by any of the standards or specifications described on the standards(5) manual page, so portable applications should not depend on it. This behavior may change in a future release.																		
%d	Day of month [1,31]; leading zero is permitted but not required.																		
%D	Date as %m/%d/%y.																		

strptime(3C)

%e	Same as %d.
%F	Equivalent to %Y-%m-%d (the ISO 8601:2000 standard date format).
%h	Same as %b.
%H	Hour (24-hour clock) [0,23]; leading zero is permitted but not required.
%I	Hour (12-hour clock) [1,12]; leading zero is permitted but not required.
%j	Day number of the year [1,366]; leading zeros are permitted but not required.
%m	Month number [1,12]; leading zero is permitted but not required.
%M	Minute [0-59]; leading zero is permitted but not required.
%n	Any white space.
%p	Locale's equivalent of either a.m. or p.m.
%r	Appropriate time representation in the 12-hour clock format with %p.
%R	Time as %H:%M.
SUSv3	
%S	Seconds [0,60]; leading zero is permitted but not required. The range of values is [00,60] rather than [00,59] to allow for the occasional leap second.
Default and other standards	
%S	Seconds [0,61]; leading zero is permitted but not required. The range of values is [00,61] rather than [00,59] to allow for the occasional leap second and even more occasional double leap second.
%t	Any white space.
%T	Time as %H:%M:%S.
%U	Week number of the year as a decimal number [0,53], with Sunday as the first day of the week; leading zero is permitted but not required.
%w	Weekday as a decimal number [0,6], with 0 representing Sunday.
%W	Week number of the year as a decimal number [0,53], with Monday as the first day of the week; leading zero is permitted but not required.
%x	Locale's appropriate date representation.
%X	Locale's appropriate time representation.
%y	Year within century. When a century is not otherwise specified, values in the range 69-99 refer to years in the twentieth century (1969 to 1999 inclusive); values in the range 00-68 refer to years in the twenty-first century (2000 to 2068 inclusive).
%Y	Year, including the century (for example, 1993).

**Modified
Conversion
Specifications**

%z Replaced by offset from UTC in ISO 8601:2000 standard format Replaced by offset from UTC in ISO 8601:2000 standard format determinable. For example, "-0430" means 4 hours 30 minutes behind UTC (west of Greenwich). If `tm_isdst` is 0, the standard time offset is used. If `tm_isdst` is greater than 0, the daylight savings time offset is used. If `tm_isdst` is negative, no characters are returned.

%Z Time zone name or no characters if no time zone exists.

Some conversion specifications can be modified by the `E` and `O` modifier characters to indicate that an alternate format or specification should be used rather than the one normally used by the unmodified specification. If the alternate format or specification does not exist in the current locale, the behavior will be as if the unmodified conversion specification were used.

%Ec Locale's alternate appropriate date and time representation.

%EC Name of the base year (era) in the locale's alternate representation.

%Ex Locale's alternate date representation.

%EX Locale's alternate time representation.

%Ey Offset from **%EC** (year only) in the locale's alternate representation.

%EY Full alternate year representation.

%Od Day of the month using the locale's alternate numeric symbols.

%Oe Same as **%Od**.

%OH Hour (24-hour clock) using the locale's alternate numeric symbols.

%OI Hour (12-hour clock) using the locale's alternate numeric symbols.

%Om Month using the locale's alternate numeric symbols.

%OM Minutes using the locale's alternate numeric symbols.

%OS Seconds using the locale's alternate numeric symbols.

%OU Week number of the year (Sunday as the first day of the week) using the locale's alternate numeric symbols.

%Ow Number of the weekday (Sunday=0) using the locale's alternate numeric symbols.

%OW Week number of the year (Monday as the first day of the week) using the locale's alternate numeric symbols.

%Oy Year (offset from **%C**) in the locale's alternate representation and using the locale's alternate numeric symbols.

strptime(3C)

General Specifications

A conversion specification that is an ordinary character is executed by scanning the next character from the buffer. If the character scanned from the buffer differs from the one comprising the specification, the specification fails, and the differing and subsequent characters remain unscanned.

A series of specifications composed of %n, %t, white-space characters or any combination is executed by scanning up to the first character that is not white space (which remains unscanned), or until no more characters can be scanned. White space is defined by `isspace(3C)`.

Any other conversion specification is executed by scanning characters until a character matching the next specification is scanned, or until no more characters can be scanned. These characters, except the one matching the next specification, are then compared to the locale values associated with the conversion specifier. If a match is found, values for the appropriate *tm* structure members are set to values corresponding to the locale information. If no match is found, `strptime()` fails and no more characters are scanned.

The month names, weekday names, era names, and alternate numeric symbols can consist of any combination of upper and lower case letters. The user can request that the input date or time specification be in a specific language by setting the `LC_TIME` category using `setlocale(3C)`.

Non-zeroing Behavior

In addition to the behavior described above by various standards, the Solaris implementation of `strptime()` provides the following extensions. These may change at any time in the future. Portable applications should not depend on these extended features:

- If `_STRPTIME_DONTZERO` is not defined, the *tm* struct is zeroed on entry and `strptime()` updates the fields of the *tm* struct associated with the specifiers in the format string.
- If `_STRPTIME_DONTZERO` is defined, `strptime()` does not zero the *tm* struct on entry. Additionally, for some specifiers, `strptime()` will use some values in the input *tm* struct to recalculate the date and re-assign the appropriate members of the *tm* struct.

The following describes extended features regardless of whether `_STRPTIME_DONTZERO` is defined or not defined:

- If %j is specified, `tm_yday` is set; if year is given, and if month and day are not given, `strptime()` calculates and sets `tm_mon`, `tm_mday`, and `tm_year`.
- If %U or %W is specified and if weekday and year are given and month and day of month are not given, `strptime()` calculates and sets `tm_mon`, `tm_mday`, `tm_wday`, and `tm_year`.

The following describes extended features when `_STRPTIME_DONTZERO` is not defined:

- If %C is specified and %y is not specified, `strptime()` assumes 0 as the year offset, then calculates the year, and assigns `tm_year`.

The following describes extended features when `_STRPTIME_DONTZERO` is defined:

- If %C is specified and %y is not specified, `strptime()` assumes the year offset of the year value of the `tm_year` member of the input `tm struct`, then calculates the year and assigns `tm_year`.
- If %j is specified and neither %y, %Y, nor %C are specified, and neither month nor day of month are specified, `strptime()` assumes the year value given by the value of the `tm_year` field of the input `tm struct`. Then, in addition to setting `tm_yday`, `strptime()` uses day-of-year and year values to calculate the month and day-of-month, and assigns `tm_month` and `tm_mday`.
- If %U or %W is specified, and if weekday and/or year are not given, and month and day of month are not given, `strptime()` will assume the weekday value and/or the year value as the value of the `tm_wday` field and/or `tm_year` field of the input `tm struct`. Then, `strptime()` will calculate the month and day-of-month and assign `tm_month`, `tm_mday`, and/or `tm_year`.
- If %p is specified and if hour is not specified, `strptime()` will reference, and if needed, update the `tm_hour` member. If the `am_pm` input is p.m. and the input `tm_hour` value is between 0 - 11, `strptime()` will add 12 hours and update `tm_hour`. If the `am_pm` input is a.m. and input `tm_hour` value is between 12 - 23, `strptime()` will subtract 12 hours and update `tm_hour`.

RETURN VALUES Upon successful completion, `strptime()` returns a pointer to the character following the last character parsed. Otherwise, a null pointer is returned.

USAGE Several “same as” formats, and the special processing of white-space characters are provided in order to ease the use of identical *format* strings for `strptime(3C)` and `strptime()`.

The `strptime()` function tries to calculate `tm_year`, `tm_mon`, and `tm_mday` when given incomplete input. This allows the `struct tm` created by `strptime()` to be passed to `mktime(3C)` to produce a `time_t` value for dates and times that are representable by a `time_t`. As an example, since `mktime()` ignores `tm_yday`, `strptime()` calculates `tm_mon` and `tm_mday` as well as filling in `tm_yday` when %j is specified without otherwise specifying a month and day within month.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe

strptime(3C)

SEE ALSO [ctime\(3C\)](#), [getdate\(3C\)](#), [isspace\(3C\)](#), [mktime\(3C\)](#), [setlocale\(3C\)](#), [strftime\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#), [standards\(5\)](#)

- NAME** strsignal – get name of signal
- SYNOPSIS** `#include <string.h>`
`char *strsignal(int sig);`
- DESCRIPTION** The `strsignal()` function maps the signal number in `sig` to a string describing the signal and returns a pointer to that string. It uses the same set of the messages as `psignal(3C)`. The returned string should not be overwritten.
- RETURN VALUES** The `strsignal()` function returns `NULL` if `sig` is not a valid signal number.
- USAGE** If the application is linked with `-lintl`, messages returned from this function are in the native language specified by the `LC_MESSAGES` locale category; see `setlocale(3C)`.
- ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `gettext(3C)`, `psignal(3C)`, `setlocale(3C)`, `str2sig(3C)`, `attributes(5)`

strtod(3C)

NAME	strtod, strtodf, strtold, atof – convert string to floating-point number
SYNOPSIS	<pre>#include <stdlib.h> double strtod(const char *restrict <i>nptr</i>, char **restrict <i>endptr</i>); float strtodf(const char *restrict <i>nptr</i>, char **restrict <i>endptr</i>); long double strtold(const char *restrict <i>nptr</i>, char **restrict <i>endptr</i>); double atof(const char *<i>str</i>);</pre>
DESCRIPTION	<p>The <code>strtod()</code>, <code>strtodf()</code>, and <code>strtold()</code> functions convert the initial portion of the string pointed to by <i>nptr</i> to double, float, and long double representation, respectively. First they decompose the input string into three parts:</p> <ol style="list-style-type: none">1. An initial, possibly empty, sequence of white-space characters (as specified by isspace(3C))2. A subject sequence interpreted as a floating-point constant or representing infinity or NaN3. A final string of one or more unrecognized characters, including the terminating null byte of the input string. <p>Then they attempt to convert the subject sequence to a floating-point number, and return the result.</p> <p>The expected form of the subject sequence is an optional plus or minus sign, then one of the following:</p> <ul style="list-style-type: none">■ A non-empty sequence of digits optionally containing a radix character, then an optional exponent part■ A 0x or 0X, then a non-empty sequence of hexadecimal digits optionally containing a radix character, then an optional binary exponent part■ One of INF or INFINITY, ignoring case■ One of NAN or NAN(<i>n-char-sequence_{opt}</i>), ignoring case in the NAN part, where: <pre>n-char-sequence: digit nondigit n-char-sequence digit n-char-sequence nondigit</pre> <p>In default mode for <code>strtod()</code>, only decimal, INF/INFINITY, and NAN/NAN(<i>n-char-sequence</i>) forms are recognized. In C99/SUSv3 mode, hexadecimal strings are also recognized.</p> <p>In default mode for <code>strtod()</code>, the <i>n-char-sequence</i> in the NAN(<i>n-char-sequence</i>) form can contain any character except ')' (right parenthesis) or '\0' (null). In C99/SUSv3 mode, the <i>n-char-sequence</i> can contain only upper and lower case letters, digits, and '_' (underscore).</p>

The `strtod()` and `strtold()` functions always function in C99/SUSv3-conformant mode.

The subject sequence is defined as the longest initial subsequence of the input string, starting with the first non-white-space character, that is of the expected form. The subject sequence contains no characters if the input string is not of the expected form.

If the subject sequence has the expected form for a floating-point number, the sequence of characters starting with the first digit or the decimal-point character (whichever occurs first) is interpreted as a floating constant of the C language, except that the radix character is used in place of a period, and that if neither an exponent part nor a radix character appears in a decimal floating-point number, or if a binary exponent part does not appear in a hexadecimal floating-point number, an exponent part of the appropriate type with value zero is assumed to follow the last digit in the string. If the subject sequence begins with a minus sign, the sequence is interpreted as negated. A character sequence `INF` or `INFINITY` is interpreted as an infinity. A character sequence `NAN` or `NAN(n-char-sequenceopt)` is interpreted as a quiet NaN. A pointer to the final string is stored in the object pointed to by `endptr`, provided that `endptr` is not a null pointer.

If the subject sequence has either the decimal or hexadecimal form, the value resulting from the conversion is rounded correctly according to the prevailing floating point rounding direction mode. The conversion also raises floating point inexact, underflow, or overflow exceptions as appropriate.

The radix character is defined in the program's locale (category `LC_NUMERIC`). In the POSIX locale, or in a locale where the radix character is not defined, the radix character defaults to a period (`.`).

If the subject sequence is empty or does not have the expected form, no conversion is performed; the value of `nptr` is stored in the object pointed to by `endptr`, provided that `endptr` is not a null pointer.

The `strtod()` function does not change the setting of `errno` if successful.

The `atof(str)` function call is equivalent to `strtod(nptr, (char **)NULL)`.

RETURN VALUES

Upon successful completion, these functions return the converted value. If no conversion could be performed, 0 is returned.

If the correct value is outside the range of representable values, `±HUGE_VAL`, `±HUGE_VALF`, or `±HUGE_VALL` is returned (according to the sign of the value), a floating point overflow exception is raised, and `errno` is set to `ERANGE`.

If the correct value would cause an underflow, the correctly rounded result (which may be normal, subnormal, or zero) is returned, a floating point underflow exception is raised, and `errno` is set to `ERANGE`.

ERRORS

These functions will fail if:

`ERANGE` The value to be returned would cause overflow or underflow

strtod(3C)

These functions may fail if:

`EINVAL` No conversion could be performed.

USAGE Since 0 is returned on error and is also a valid return on success, an application wishing to check for error situations should set `errno` to 0, then call `strtod()`, `strtof()`, or `strtold()`, then check `errno`.

The changes to `strtod()` introduced by the ISO/IEC 9899:1999 standard can alter the behavior of well-formed applications complying with the ISO/IEC 9899:1990 standard and thus earlier versions of IEEE Std 1003.1-200x. One such example would be:

```
int
what_kind_of_number (char *s)
{
    char *endp;
    double d;
    long l;
    d = strtod(s, &endp);
    if (s != endp && *endp == '\\0')
        printf("It's a float with value %g\n", d);
    else
    {
        l = strtol(s, &endp, 0);
        if (s != endp && *endp == '\\0')
            printf("It's an integer with value %ld\n", l);
        else
            return 1;
    }
    return 0;
}
```

If the function is called with:

```
what_kind_of_number ("0x10")
```

an ISO/IEC 9899:1990 standard-compliant library will result in the function printing:

```
It's an integer with value 16
```

With the ISO/IEC 9899:1999 standard, the result is:

```
It's a float with value 16
```

The change in behavior is due to the inclusion of floating-point numbers in hexadecimal notation without requiring that either a decimal point or the binary exponent be present.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

SEE ALSO [isspace\(3C\)](#), [localeconv\(3C\)](#), [scanf\(3C\)](#), [setlocale\(3C\)](#), [strtol\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NOTES The `strtod()` and `atof()` functions can be used safely in multithreaded applications, as long as [setlocale\(3C\)](#) is not called to change the locale.

The DESCRIPTION and RETURN VALUES sections above are very similar to the wording used by the Single UNIX Specification version 2 (SUSv2) and the 1989 C Standard to describe the behavior of the `strtod()` function. Since some users have reported that they find the description confusing, the following notes might be helpful.

1. The `strtod()` function does not modify the string pointed to by *str* and does not `malloc()` space to hold the decomposed portions of the input string.
2. If *endptr* is not `(char **)NULL`, `strtod()` will set the pointer pointed to by *endptr* to the first byte of the "final string of unrecognized characters". (If all input characters were processed, the pointer pointed to by *endptr* will be set to point to the null character at the end of the input string.)
3. If `strtod()` returns 0.0, one of the following occurred:
 - a. The "subject sequence" was not an empty string, but evaluated to 0.0. (In this case, `errno` will be left unchanged.)
 - b. The "subject sequence" was an empty string. In this case, `errno` will be left unchanged. (The Single UNIX Specification version 2 allows `errno` to be set to `EINVAL` or to be left unchanged. The C Standard does not specify any specific behavior in this case.)
 - c. The "subject sequence" specified a numeric value whose conversion resulted in a floating point underflow. In this case, an underflow exception is raised and `errno` is set to `ERANGE`.

Note that the standards do not require that implementations distinguish between these three cases. An application can determine case (b) by making sure that there are no leading white-space characters in the string pointed to by *str* and giving `strtod()` an *endptr* that is not `(char **)NULL`. If *endptr* points to the first character of *str* when `strtod()` returns, you have detected case (b). Case (c) can be detected by examining the underflow flag or by looking for a non-zero digit before the exponent part of the "subject sequence". Note, however, that the decimal-point character is locale-dependent.

4. If `strtod()` returns `+HUGE_VAL` or `-HUGE_VAL`, one of the following occurred:

strtod(3C)

- a. If `+HUGE_VAL` is returned and `errno` is set to `ERANGE`, a floating point overflow occurred while processing a positive value, causing a floating point overflow exception to be raised.
- b. If `-HUGE_VAL` is returned and `errno` is set to `ERANGE`, a floating point overflow occurred while processing a negative value, causing a floating point overflow exception to be raised.
- c. If `strtod()` does not set `errno` to `ERANGE`, the value specified by the "subject string" converted to `+HUGE_VAL` or `-HUGE_VAL`, respectively.

Note that if `errno` is set to `ERANGE` when `strtod()` is called, case (c) can be distinguished from cases (a) and (b) by examining either `ERANGE` or the overflow flag.

NAME	strtoimax, strtoumax – convert string to integer type						
SYNOPSIS	<pre>#include <inttypes.h> intmax_t strtoimax(const char *restrict <i>nptr</i>, char **restrict <i>endptr</i>, int <i>base</i>); uintmax_t strtoumax(const char *restrict <i>nptr</i>, char **restrict <i>endptr</i>, int <i>base</i>);</pre>						
DESCRIPTION	These functions are equivalent to the <code>strtol()</code> , <code>strtoll()</code> , <code>strtoul()</code> , and <code>strtoull()</code> functions, except that the initial portion of the string is converted to <code>intmax_t</code> and <code>uintmax_t</code> representation, respectively.						
RETURN VALUES	<p>These functions return the converted value, if any.</p> <p>If no conversion could be performed, 0 is returned.</p> <p>If the correct value is outside the range of representable values, <code>{INTMAX_MAX}</code>, <code>{INTMAX_MIN}</code>, or <code>{UINTMAX_MAX}</code> is returned (according to the return type and sign of the value, if any), and <code>errno</code> is set to <code>ERANGE</code>.</p>						
ERRORS	<p>These functions will fail if:</p> <p><code>ERANGE</code> The value to be returned is not representable.</p> <p>These functions may fail if:</p> <p><code>EINVAL</code> The value of <i>base</i> is not supported.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th style="text-align: center;">ATTRIBUTE TYPE</th> <th style="text-align: center;">ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>strtol(3C)</code> , <code>strtoul(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

strtol(3C)

NAME	strtol, strtoll, atol, atoll, atoi, lltostr, ulltostr – string conversion routines
SYNOPSIS	<pre>#include <stdlib.h> long strtol(const char *restrict <i>str</i>, char **restrict <i>endptr</i>, int <i>base</i>); long long strtoll(const char *restrict <i>str</i>, char **restrict <i>endptr</i>, int <i>base</i>); long atol(const char *<i>str</i>); long long atoll(const char *<i>str</i>); int atoi(const char *<i>str</i>); char *lltostr(long long <i>value</i>, char *<i>endptr</i>); char *ulltostr(unsigned long long <i>value</i>, char *<i>endptr</i>);</pre>
DESCRIPTION	
strtol() and strtoll()	<p>The <code>strtol()</code> function converts the initial portion of the string pointed to by <i>str</i> to a type long int representation.</p> <p>The <code>strtoll()</code> function converts the initial portion of the string pointed to by <i>str</i> to a type long long representation.</p> <p>Both functions first decompose the input string into three parts: an initial, possibly empty, sequence of white-space characters (as specified by <code>isspace(3C)</code>); a subject sequence interpreted as an integer represented in some radix determined by the value of <i>base</i>; and a final string of one or more unrecognized characters, including the terminating null byte of the input string. They then attempt to convert the subject sequence to an integer and return the result.</p> <p>If the value of <i>base</i> is 0, the expected form of the subject sequence is that of a decimal constant, octal constant or hexadecimal constant, any of which may be preceded by a + or – sign. A decimal constant begins with a non-zero digit, and consists of a sequence of decimal digits. An octal constant consists of the prefix 0 optionally followed by a sequence of the digits 0 to 7 only. A hexadecimal constant consists of the prefix 0x or 0X followed by a sequence of the decimal digits and letters a (or A) to f (or F) with values 10 to 15 respectively.</p> <p>If the value of <i>base</i> is between 2 and 36, the expected form of the subject sequence is a sequence of letters and digits representing an integer with the radix specified by <i>base</i>, optionally preceded by a + or – sign. The letters from a (or A) to z (or Z) inclusive are ascribed the values 10 to 35; only letters whose ascribed values are less than that of <i>base</i> are permitted. If the value of <i>base</i> is 16, the characters 0x or 0X may optionally precede the sequence of letters and digits, following the sign if present.</p>

The subject sequence is defined as the longest initial subsequence of the input string, starting with the first non-white-space character, that is of the expected form. The subject sequence contains no characters if the input string is empty or consists entirely of white-space characters, or if the first non-white-space character is other than a sign or a permissible letter or digit.

If the subject sequence has the expected form and the value of *base* is 0, the sequence of characters starting with the first digit is interpreted as an integer constant. If the subject sequence has the expected form and the value of *base* is between 2 and 36, it is used as the base for conversion, ascribing to each letter its value as given above. If the subject sequence begins with a minus sign, the value resulting from the conversion is negated. A pointer to the final string is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

In other than the POSIX locale, additional implementation-dependent subject sequence forms may be accepted.

If the subject sequence is empty or does not have the expected form, no conversion is performed; the value of *str* is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

**atol(), atoll()
and atoi()**

Except for behavior on error, `atol()` is equivalent to: `strtoul(str, (char **)NULL, 10)`.

Except for behavior on error, `atoll()` is equivalent to: `strtoull(str, (char **)NULL, 10)`.

Except for behavior on error, `atoi()` is equivalent to: `(int) strtoul(str, (char **)NULL, 10)`.

If the value cannot be represented, the behavior is undefined.

**lltostr() and
ulltostr()**

The `lltostr()` function returns a pointer to the string represented by the long long *value*. The *endptr* argument is assumed to point to the byte following a storage area into which the decimal representation of *value* is to be placed as a string. The `lltostr()` function converts *value* to decimal and produces the string, and returns a pointer to the beginning of the string. No leading zeros are produced, and no terminating null is produced. The low-order digit of the result always occupies memory position *endptr*-1. The behavior of `lltostr()` is undefined if *value* is negative. A single zero digit is produced if *value* is 0.

The `ulltostr()` function is similar to `lltostr()` except that *value* is an unsigned long long.

RETURN VALUES

Upon successful completion, `strtoul()`, `strtoull()`, `atol()`, `atoll()`, and `atoi()` return the converted value, if any. If no conversion could be performed, `strtoul()` and `strtoull()` return 0 and `errno` may be set to `EINVAL`.

If the correct value is outside the range of representable values, `strtoul()` returns `LONG_MAX` or `LONG_MIN` and `strtoull()` returns `LLONG_MAX` or `LLONG_MIN` (according to the sign of the value), and `errno` is set to `ERANGE`.

strtol(3C)

Upon successful completion, `lltostr()` and `ulltostr()` return a pointer to the converted string.

ERRORS The `strtol()` and `strtoll()` functions will fail if:

`ERANGE` The value to be returned is not representable. The `strtol()` and `strtoll()` functions may fail if:

`EINVAL` The value of *base* is not supported.

USAGE Because 0, `LONG_MIN`, `LONG_MAX`, `LLONG_MIN`, and `LLONG_MAX` are returned on error and are also valid returns on success, an application wishing to check for error situations should set `errno` to 0, call the function, then check `errno` and if it is non-zero, assume an error has occurred.

The `strtol()` function no longer accepts values greater than `LONG_MAX` or `LLONG_MAX` as valid input. Use [strtoul\(3C\)](#) instead.

Calls to `atoi()` and `atol()` might be faster than corresponding calls to `strtol()`, and calls to `atoll()` might be faster than corresponding calls to `strtoll()`. However, applications should not use the `atoi()`, `atol()`, or `atoll()` functions unless they know the value represented by the argument will be in range for the corresponding result type.

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	See below.
MT-Level	MT-Safe

The `strtol()`, `strtoll()`, `atol()`, `atoll()`, and `atoi()` functions are Standard.

SEE ALSO [isalpha\(3C\)](#), [isspace\(3C\)](#), [scanf\(3C\)](#), [strtod\(3C\)](#), [strtoul\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	strtol, strtoll – convert string to unsigned long
SYNOPSIS	<pre>#include <stdlib.h> unsigned long strtol(const char *restrict <i>str</i>, char **restrict <i>endptr</i>, int <i>base</i>); unsigned long long strtoll(const char *restrict <i>str</i>, char **restrict <i>endptr</i>, int <i>base</i>);</pre>
DESCRIPTION	<p>The <code>strtol()</code> function converts the initial portion of the string pointed to by <i>str</i> to a type unsigned long int representation. First it decomposes the input string into three parts: an initial, possibly empty, sequence of white-space characters (as specified by <code>isspace(3C)</code>); a subject sequence interpreted as an integer represented in some radix determined by the value of <i>base</i>; and a final string of one or more unrecognised characters, including the terminating null byte of the input string. Then it attempts to convert the subject sequence to an unsigned integer, and returns the result.</p> <p>If the value of <i>base</i> is 0, the expected form of the subject sequence is that of a decimal constant, octal constant or hexadecimal constant, any of which may be preceded by a + or – sign. A decimal constant begins with a non-zero digit, and consists of a sequence of decimal digits. An octal constant consists of the prefix 0 optionally followed by a sequence of the digits 0 to 7 only. A hexadecimal constant consists of the prefix 0x or 0X followed by a sequence of the decimal digits and letters a (or A) to f (or F) with values 10 to 15 respectively.</p> <p>If the value of <i>base</i> is between 2 and 36, the expected form of the subject sequence is a sequence of letters and digits representing an integer with the radix specified by <i>base</i>, optionally preceded by a + or – sign. The letters from a (or A) to z (or Z) inclusive are ascribed the values 10 to 35; only letters whose ascribed values are less than that of <i>base</i> are permitted. If the value of <i>base</i> is 16, the characters 0x or 0X may optionally precede the sequence of letters and digits, following the sign if present.</p> <p>The subject sequence is defined as the longest initial subsequence of the input string, starting with the first non-white-space character, that is of the expected form. The subject sequence contains no characters if the input string is empty or consists entirely of white-space characters, or if the first non-white-space character is other than a sign or a permissible letter or digit.</p> <p>If the subject sequence has the expected form and the value of <i>base</i> is 0, the sequence of characters starting with the first digit is interpreted as an integer constant. If the subject sequence has the expected form and the value of <i>base</i> is between 2 and 36, it is used as the base for conversion, ascribing to each letter its value as given above. If the subject sequence begins with a minus sign, the value resulting from the conversion is negated. A pointer to the final string is stored in the object pointed to by <i>endptr</i>, provided that <i>endptr</i> is not a null pointer.</p> <p>In other than the POSIX locale, additional implementation-dependent subject sequence forms may be accepted.</p>

strtoul(3C)

If the subject sequence is empty or does not have the expected form, no conversion is performed; the value of *str* is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

The `strtoull()` function is identical to `strtoul()` except that it returns the value represented by *str* as an unsigned long long.

RETURN VALUES

Upon successful completion `strtoul()` returns the converted value, if any. If no conversion could be performed, 0 is returned and `errno` may be set to `EINVAL`. If the correct value is outside the range of representable values, `ULONG_MAX` is returned and `errno` is set to `ERANGE`.

ERRORS

The `strtoul()` function will fail if:

`EINVAL` The value of *base* is not supported.

`ERANGE` The value to be returned is not representable.

The `strtoul()` function may fail if:

`EINVAL` No conversion could be performed.

USAGE

Because 0 and `ULONG_MAX` are returned on error and are also valid returns on success, an application wishing to check for error situations should set `errno` to 0, then call `strtoul()`, then check `errno` and if it is non-zero, assume an error has occurred.

Unlike `strtod(3C)` and `strtol(3C)`, `strtoul()` must always return a non-negative number; so, using the return value of `strtoul()` for out-of-range numbers with `strtoul()` could cause more severe problems than just loss of precision if those numbers can ever be negative.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO

`isalpha(3C)`, `isspace(3C)`, `scanf(3C)`, `strtod(3C)`, `strtol(3C)`, `attributes(5)`, `standards(5)`

NAME	strtoks, wstoktr – code conversion for Process Code and File Code
SYNOPSIS	<pre>#include <widec.h> wchar_t *strtoks(wchar_t *dst, const char *src); char *wstoktr(char *dst, const wchar_t *src);</pre>
DESCRIPTION	<p>The <code>strtoks()</code> and <code>wstoktr()</code> functions convert strings back and forth between File Code representation and Process Code.</p> <p>The <code>strtoks()</code> function takes a character string <i>src</i>, converts it to a Process Code string, terminated by a Process Code null, and places the result into <i>dst</i>.</p> <p>The <code>wstoktr()</code> function takes the Process Code string pointed to by <i>src</i>, converts it to a character string, and places the result into <i>dst</i>.</p>
RETURN VALUES	<p>The <code>strtoks()</code> function returns the Process Code string if it completes successfully. Otherwise, a null pointer will be returned and <code>errno</code> will be set to <code>EILSEQ</code>.</p> <p>The <code>wstoktr()</code> function returns the File Code string if it completes successfully. Otherwise, a null pointer will be returned and <code>errno</code> will be set to <code>EILSEQ</code>.</p>
SEE ALSO	wstring(3C)

strxfrm(3C)

NAME	strxfrm – string transformation
SYNOPSIS	<pre>#include <string.h> size_t strxfrm(char *restrict s1, const char *restrict s2, size_t n);</pre>
DESCRIPTION	<p>The <code>strxfrm()</code> function transforms the string pointed to by <code>s2</code> and places the resulting string into the array pointed to by <code>s1</code>. The transformation is such that if <code>strcmp(3C)</code> is applied to two transformed strings, it returns a value greater than, equal to or less than 0, corresponding to the result of <code>strcoll(3C)</code> applied to the same two original strings. No more than <code>n</code> bytes are placed into the resulting array pointed to by <code>s1</code>, including the terminating null byte. If <code>n</code> is 0, <code>s1</code> is permitted to be a null pointer. If copying takes place between objects that overlap, the behavior is undefined.</p> <p>The <code>strxfrm()</code> function does not change the setting of <code>errno</code> if successful.</p> <p>Since no return value is reserved to indicate an error, an application wishing to check for error situations should set <code>errno</code> to 0, then call <code>strxfrm()</code>, then check <code>errno</code>.</p>
RETURN VALUES	<p>Upon successful completion, <code>strxfrm()</code> returns the length of the transformed string (not including the terminating null byte). If the value returned is <code>n</code> or more, the contents of the array pointed to by <code>s1</code> are indeterminate.</p> <p>On error, <code>strxfrm()</code> may set <code>errno</code> but no return value is reserved to indicate the error.</p>
USAGE	<p>The transformation function is such that two transformed strings can be ordered by <code>strcmp(3C)</code> as appropriate to collating sequence information in the program's locale (category <code>LC_COLLATE</code>).</p> <p>The fact that when <code>n</code> is 0, <code>s1</code> is permitted to be a null pointer, is useful to determine the size of the <code>s1</code> array prior to making the transformation.</p>
EXAMPLES	<p>EXAMPLE 1 A sample of using the <code>strxfrm()</code> function.</p> <p>The value of the following expression is the size of the array needed to hold the transformation of the string pointed to by <code>s</code>.</p> <pre>1 + strxfrm(NULL, s, 0);</pre>
FILES	<pre>/usr/lib/locale/locale/locale.so.*</pre> LC_COLLATE database for <i>locale</i>
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe with exceptions

The `strxfrm()` function can be used safely in a multithreaded application, as long as `setlocale(3C)` is not being called to change the locale.

SEE ALSO `localedef(1)`, `setlocale(3C)`, `strcmp(3C)`, `strcoll(3C)`, `wscoll(3C)`, `attributes(5)`, `environ(5)`, `standards(5)`

swab(3C)

NAME	swab – swap bytes						
SYNOPSIS	<pre>#include <stdlib.h> void swab(const char *src, char *dest, ssize_t nbytes);</pre>						
XPG4, SUS, SUSv2, SUSv3	<pre>#include <unistd.h> void swab(const void *restrict src, void *restrict dest, ssize_t nbytes);</pre>						
DESCRIPTION	<p>The <code>swab()</code> function copies <i>nbytes</i> bytes, which are pointed to by <i>src</i>, to the object pointed to by <i>dest</i>, exchanging adjacent bytes. The <i>nbytes</i> argument should be even. If <i>nbytes</i> is odd <code>swab()</code> copies and exchanges <i>nbytes</i>–1 bytes and the disposition of the last byte is unspecified. If copying takes place between objects that overlap, the behavior is undefined. If <i>nbytes</i> is negative, <code>swab()</code> does nothing.</p>						
ERRORS	No errors are defined.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>attributes(5)</code> , <code>standards(5)</code>						

NAME	sync_instruction_memory – make modified instructions executable				
SYNOPSIS	void sync_instruction_memory (caddr_t <i>addr</i> , int <i>len</i>);				
DESCRIPTION	<p>The <code>sync_instruction_memory()</code> function performs whatever steps are required to make instructions modified by a program executable.</p> <p>Some processor architectures, including some SPARC processors, have separate and independent instruction and data caches which are not kept consistent by hardware. For example, if the instruction cache contains an instruction from some address and the program then stores a new instruction at that address, the new instruction may not be immediately visible to the instruction fetch mechanism. Software must explicitly invalidate the instruction cache entries for new or changed mappings of pages that might contain executable instructions. The <code>sync_instruction_memory()</code> function performs this function, and/or any other functions needed to make modified instructions between <i>addr</i> and <i>addr+len</i> visible. A program should call <code>sync_instruction_memory()</code> after modifying instructions and before executing them.</p> <p>On processors with unified caches (one cache for both instructions and data) and pipelines which are flushed by a branch instruction, such as the x86 architecture, the function may do nothing and just return.</p> <p>The changes are immediately visible to the thread calling <code>sync_instruction_memory()</code> when the call returns, even if the thread should migrate to another processor during or after the call. The changes become visible to other threads in the same manner that stores do; that is, they eventually become visible, but the latency is implementation-dependent.</p> <p>The result of executing <code>sync_instruction_memory()</code> are unpredictable if <i>addr</i> through <i>addr+len-1</i> are not valid for the address space of the program making the call.</p>				
RETURN VALUES	No values are returned.				
ATTRIBUTES	See <code>attributes(5)</code> 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>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>attributes(5)</code>				

syscall(3UCB)

NAME	syscall – indirect system call
SYNOPSIS	<pre>/usr/ucb/cc [flag ...] file ... #include <sys/syscall.h> int syscall (number, arg, ...);</pre>
DESCRIPTION	syscall() performs the function whose assembly language interface has the specified <i>number</i> , and arguments <i>arg</i> ... Symbolic constants for functions can be found in the header <sys/syscall.h>.
RETURN VALUES	On error syscall() returns -1 and sets the external variable errno (see intro(2)).
FILES	<sys/syscall.h>
SEE ALSO	intro(2), pipe(2)
NOTES	Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.
WARNINGS	<p>There is no way to use syscall() to call functions such as pipe(2) which return values that do not fit into one hardware register.</p> <p>Since many system calls are implemented as library wrappers around traps to the kernel, these calls may not behave as documented when called from syscall(), which bypasses these wrappers. For these reasons, using syscall() is not recommended.</p>

NAME	sysconf – get configurable system variables																																										
SYNOPSIS	<pre>#include <unistd.h> long sysconf(int <i>name</i>);</pre>																																										
DESCRIPTION	<p>The <code>sysconf()</code> function provides a method for an application to determine the current value of a configurable system limit or option (variable).</p> <p>The <i>name</i> argument represents the system variable to be queried. The following table lists the minimal set of system variables from <code><limits.h></code> and <code><unistd.h></code> that can be returned by <code>sysconf()</code> and the symbolic constants defined in <code><unistd.h></code> that are the corresponding values used for <i>name</i> on the SPARC and x86 platforms.</p> <table border="1"> <thead> <tr> <th>Name</th> <th>Return Value</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td><code>_SC_2_C_BIND</code></td> <td><code>_POSIX2_C_BIND</code></td> <td>Supports the C language binding option</td> </tr> <tr> <td><code>_SC_2_C_DEV</code></td> <td><code>_POSIX2_C_DEV</code></td> <td>Supports the C language development utilities option</td> </tr> <tr> <td><code>_SC_2_C_VERSION</code></td> <td><code>_POSIX2_C_VERSION</code></td> <td>Integer value indicates version of ISO POSIX-2 standard (Commands)</td> </tr> <tr> <td><code>_SC_2_CHAR_TERM</code></td> <td><code>_POSIX2_CHAR_TERM</code></td> <td>Supports at least one terminal</td> </tr> <tr> <td><code>_SC_2_FORT_DEV</code></td> <td><code>_POSIX2_FORT_DEV</code></td> <td>Supports FORTRAN Development Utilities Option</td> </tr> <tr> <td><code>_SC_2_FORT_RUN</code></td> <td><code>_POSIX2_FORT_RUN</code></td> <td>Supports FORTRAN Run-time Utilities Option</td> </tr> <tr> <td><code>_SC_2_LOCALEDEF</code></td> <td><code>_POSIX2_LOCALEDEF</code></td> <td>Supports creation of locales by the localedef utility</td> </tr> <tr> <td><code>_SC_2_SW_DEV</code></td> <td><code>_POSIX2_SW_DEV</code></td> <td>Supports Software Development Utility Option</td> </tr> <tr> <td><code>_SC_2_UPE</code></td> <td><code>_POSIX2_UPE</code></td> <td>Supports User Portability Utilities Option</td> </tr> <tr> <td><code>_SC_2_VERSION</code></td> <td><code>_POSIX2_VERSION</code></td> <td>Integer value indicates version of ISO POSIX-2 standard (C language binding)</td> </tr> <tr> <td><code>_SC_AIO_LISTIO_MAX</code></td> <td><code>AIO_LISTIO_MAX</code></td> <td>Max number of I/O operations in a single list I/O call supported</td> </tr> <tr> <td><code>_SC_AIO_MAX</code></td> <td><code>AIO_MAX</code></td> <td>Max number of outstanding asynchronous I/O operations supported</td> </tr> <tr> <td><code>_SC_AIO_PRIO_DELTA_MAX</code></td> <td><code>AIO_PRIO_DELTA_MAX</code></td> <td>Max amount by which process can decrease</td> </tr> </tbody> </table>	Name	Return Value	Meaning	<code>_SC_2_C_BIND</code>	<code>_POSIX2_C_BIND</code>	Supports the C language binding option	<code>_SC_2_C_DEV</code>	<code>_POSIX2_C_DEV</code>	Supports the C language development utilities option	<code>_SC_2_C_VERSION</code>	<code>_POSIX2_C_VERSION</code>	Integer value indicates version of ISO POSIX-2 standard (Commands)	<code>_SC_2_CHAR_TERM</code>	<code>_POSIX2_CHAR_TERM</code>	Supports at least one terminal	<code>_SC_2_FORT_DEV</code>	<code>_POSIX2_FORT_DEV</code>	Supports FORTRAN Development Utilities Option	<code>_SC_2_FORT_RUN</code>	<code>_POSIX2_FORT_RUN</code>	Supports FORTRAN Run-time Utilities Option	<code>_SC_2_LOCALEDEF</code>	<code>_POSIX2_LOCALEDEF</code>	Supports creation of locales by the localedef utility	<code>_SC_2_SW_DEV</code>	<code>_POSIX2_SW_DEV</code>	Supports Software Development Utility Option	<code>_SC_2_UPE</code>	<code>_POSIX2_UPE</code>	Supports User Portability Utilities Option	<code>_SC_2_VERSION</code>	<code>_POSIX2_VERSION</code>	Integer value indicates version of ISO POSIX-2 standard (C language binding)	<code>_SC_AIO_LISTIO_MAX</code>	<code>AIO_LISTIO_MAX</code>	Max number of I/O operations in a single list I/O call supported	<code>_SC_AIO_MAX</code>	<code>AIO_MAX</code>	Max number of outstanding asynchronous I/O operations supported	<code>_SC_AIO_PRIO_DELTA_MAX</code>	<code>AIO_PRIO_DELTA_MAX</code>	Max amount by which process can decrease
Name	Return Value	Meaning																																									
<code>_SC_2_C_BIND</code>	<code>_POSIX2_C_BIND</code>	Supports the C language binding option																																									
<code>_SC_2_C_DEV</code>	<code>_POSIX2_C_DEV</code>	Supports the C language development utilities option																																									
<code>_SC_2_C_VERSION</code>	<code>_POSIX2_C_VERSION</code>	Integer value indicates version of ISO POSIX-2 standard (Commands)																																									
<code>_SC_2_CHAR_TERM</code>	<code>_POSIX2_CHAR_TERM</code>	Supports at least one terminal																																									
<code>_SC_2_FORT_DEV</code>	<code>_POSIX2_FORT_DEV</code>	Supports FORTRAN Development Utilities Option																																									
<code>_SC_2_FORT_RUN</code>	<code>_POSIX2_FORT_RUN</code>	Supports FORTRAN Run-time Utilities Option																																									
<code>_SC_2_LOCALEDEF</code>	<code>_POSIX2_LOCALEDEF</code>	Supports creation of locales by the localedef utility																																									
<code>_SC_2_SW_DEV</code>	<code>_POSIX2_SW_DEV</code>	Supports Software Development Utility Option																																									
<code>_SC_2_UPE</code>	<code>_POSIX2_UPE</code>	Supports User Portability Utilities Option																																									
<code>_SC_2_VERSION</code>	<code>_POSIX2_VERSION</code>	Integer value indicates version of ISO POSIX-2 standard (C language binding)																																									
<code>_SC_AIO_LISTIO_MAX</code>	<code>AIO_LISTIO_MAX</code>	Max number of I/O operations in a single list I/O call supported																																									
<code>_SC_AIO_MAX</code>	<code>AIO_MAX</code>	Max number of outstanding asynchronous I/O operations supported																																									
<code>_SC_AIO_PRIO_DELTA_MAX</code>	<code>AIO_PRIO_DELTA_MAX</code>	Max amount by which process can decrease																																									

sysconf(3C)

		its asynchronous I/O priority level from its own scheduling priority
_SC_ARG_MAX	ARG_MAX	Max size of argv[] plus envp[]
_SC_ASYNCHRONOUS_IO	_POSIX_ASYNCHRONOUS_IO	Supports Asynchronous I/O
_SC_ATEXIT_MAX	ATEXIT_MAX	Max number of functions that can be registered with atexit()
_SC_AVPHYS_PAGES		Number of physical memory pages not currently in use by system
_SC_BARRIERS	_POSIX_BARRIERS	Supports Barriers option
_SC_BC_BASE_MAX	BC_BASE_MAX	Maximum obase values allowed by bc
_SC_BC_DIM_MAX	BC_DIM_MAX	Max number of elements permitted in array by bc
_SC_BC_SCALE_MAX	BC_SCALE_MAX	Max scale value allowed by bc
_SC_BC_STRING_MAX	BC_STRING_MAX	Max length of string constant allowed by bc
_SC_CHILD_MAX	CHILD_MAX	Max processes allowed to a UID
_SC_CLK_TCK	CLK_TCK	Ticks per second (clock_t)
_SC_CLOCK_SELECTION	_POSIX_CLOCK_SELECTION	Supports Clock Selection option
_SC_COLL_WEIGHTS_MAX	COLL_WEIGHTS_MAX	Max number of weights that can be assigned to entry of the LC_COLLATE order keyword in locale definition file
_SC_CPUID_MAX		Max possible processor ID
_SC_DELAYTIMER_MAX	DELAYTIMER_MAX	Max number of timer expiration overruns
_SC_EXPR_NEST_MAX	EXPR_NEST_MAX	Max number of parentheses by expr
_SC_FSYNC	_POSIX_FSYNC	Supports File Synchronization
_SC_GETGR_R_SIZE_MAX	NSS_BUFLN_GROUP	Max size of group entry buffer
_SC_GETPW_R_SIZE_MAX	NSS_BUFLN_PASSWD	Max size of password entry buffer
_SC_HOST_NAME_MAX	_POSIX_HOST_NAME_MAX	Maximum length of a host name (excluding terminating null)
_SC_IOV_MAX	IOV_MAX	Max number of iovec structures available to one process for

		use with readv() and writev() Job control supported?
_SC_JOB_CONTROL	_POSIX_JOB_CONTROL	
_SC_LINE_MAX	LINE_MAX	Max length of input line
_SC_LOGIN_NAME_MAX	LOGNAME_MAX + 1	Max length of login name
_SC_LOGNAME_MAX	LOGNAME_MAX	
_SC_MAPPED_FILES	_POSIX_MAPPED_FILES	Supports Memory Mapped Files
_SC_MAXPID		Max pid value
_SC_MEMLOCK	_POSIX_MEMLOCK	Supports Process Memory Locking
_SC_MEMLOCK_RANGE	_POSIX_MEMLOCK_RANGE	Supports Range Memory Locking
_SC_MEMORY_PROTECTION	_POSIX_MEMORY_PROTECTION	Supports Memory Protection
_SC_MESSAGE_PASSING	_POSIX_MESSAGE_PASSING	Supports Message Passing
_SC_MONOTONIC_CLOCK	_POSIX_MONOTONIC_CLOCK	Supports Monotonic Clock option
_SC_MQ_OPEN_MAX	MQ_OPEN_MAX	Max number of open message queues a process can hold
_SC_MQ_PRIO_MAX	MQ_PRIO_MAX	Max number of message priorities supported
_SC_NGROUPS_MAX	NGROUPS_MAX	Max simultaneous groups to which one can belong
_SC_NPROCESSORS_CONF		Number of processors configured
_SC_NPROCESSORS_MAX		Max number of processors supported by platform
_SC_NPROCESSORS_ONLN		Number of processors online
_SC_OPEN_MAX	OPEN_MAX	Max open files per process
_SC_PAGESIZE	PAGESIZE	System memory page size
_SC_PAGE_SIZE	PAGESIZE	Same as _SC_PAGESIZE
_SC_PASS_MAX	PASS_MAX	Max number of significant bytes in a password
_SC_PHYS_PAGES		Total number of pages of physical memory in system
_SC_PRIORITIZED_IO	_POSIX_PRIORITIZED_IO	Supports Prioritized I/O
_SC_PRIORITY_SCHEDULING	_POSIX_PRIORITY_SCHEDULING	Supports Process Scheduling
_SC_RAW_SOCKETS	_POSIX_RAW_SOCKETS	Supports Raw Sockets option
_SC_RE_DUP_MAX	RE_DUP_MAX	Max number of repeated occurrences of a regular

sysconf(3C)

		expression permitted when using interval notation $\{m,n\}$
<code>_SC_READER_WRITER_LOCKS</code>	<code>_POSIX_READER_WRITER_LOCKS</code>	Supports IPV6 option
<code>_SC_REALTIME_SIGNALS</code>	<code>_POSIX_REALTIME_SIGNALS</code>	Supports Realtime Signals
<code>_SC_REGEX</code>	<code>_POSIX_REGEX</code>	Supports Regular Expression Handling option
<code>_SC_RTSIG_MAX</code>	<code>RTSIG_MAX</code>	Max number of realtime signals reserved for application use
<code>_SC_SAVED_IDS</code>	<code>_POSIX_SAVED_IDS</code>	Saved IDs (seteuid()) supported?
<code>_SC_SEM_NSEMS_MAX</code>	<code>SEM_NSEMS_MAX</code>	Max number of POSIX semaphores a process can have
<code>_SC_SEM_VALUE_MAX</code>	<code>SEM_VALUE_MAX</code>	Max value a POSIX semaphore can have
<code>_SC_SEMAPHORES</code>	<code>_POSIX_SEMAPHORES</code>	Supports Semaphores
<code>_SC_SHARED_MEMORY_OBJECTS</code>	<code>_POSIX_SHARED_MEMORY_OBJECTS</code>	Supports Shared Memory Objects
<code>_SC_SHELL</code>	<code>_POSIX_SHELL</code>	Supports POSIX shell
<code>_SC_SIGQUEUE_MAX</code>	<code>SIGQUEUE_MAX</code>	Max number of queued signals that a process can send and have pending at receiver(s) at a time
<code>_SC_SPAWN</code>	<code>_POSIX_SPAWN</code>	Supports Spawn option
<code>_SC_SPIN_LOCKS</code>	<code>_POSIX_SPIN_LOCKS</code>	Supports Spin Locks option
<code>_SC_STACK_PROT</code>		Default stack protection
<code>_SC_STREAM_MAX</code>	<code>STREAM_MAX</code>	Number of streams one process can have open at a time
<code>_SC_SYMLINK_MAX</code>	<code>_POSIX_SYMLINK_MAX</code>	Max number of symbolic links that can be reliably traversed in the resolution of a pathname in the absence of a loop
<code>_SC_SYNCHRONIZED_IO</code>	<code>_POSIX_SYNCHRONIZED_IO</code>	Supports Synchronized I/O
<code>_SC_THREAD_ATTR_STACKADDR</code>	<code>_POSIX_THREAD_ATTR_STACKADDR</code>	Supports Thread Stack Address Attribute option
<code>_SC_THREAD_ATTR_STACKSIZE</code>	<code>_POSIX_THREAD_ATTR_STACKSIZE</code>	Supports Thread Stack Size Attribute option
<code>_SC_THREAD_DESTRUCTOR_ITERATIONS</code>	<code>PTHREAD_DESTRUCTOR_ITERATIONS</code>	Number attempts made to destroy thread-specific data on thread exit
<code>_SC_THREAD_KEYS_MAX</code>	<code>PTHREAD_KEYS_MAX</code>	Max number of data

<code>_SC_THREAD_PRIO_INHERIT</code>	<code>_POSIX_THREAD_PRIO_INHERIT</code>	keys per process Supports Priority Inheritance option
<code>_SC_THREAD_PRIO_PROTECT</code>	<code>_POSIX_THREAD_PRIO_PROTECT</code>	Supports Priority Protection option
<code>_SC_THREAD_PRIORITY_SCHEDULING</code>	<code>_POSIX_THREAD_PRIORITY_SCHEDULING</code>	Supports Thread Execution Scheduling option
<code>_SC_THREAD_PROCESS_SHARED</code>	<code>_POSIX_THREAD_PROCESS_SHARED</code>	Supports Process-Shared Synchronization option
<code>_SC_THREAD_SAFE_FUNCTIONS</code>	<code>_POSIX_THREAD_SAFE_FUNCTIONS</code>	Supports Thread-Safe Functions option
<code>_SC_THREAD_STACK_MIN</code>	<code>PTHREAD_STACK_MIN</code>	Min byte size of thread stack storage
<code>_SC_THREAD_THREADS_MAX</code>	<code>PTHREAD_THREADS_MAX</code>	Max number of threads per process
<code>_SC_THREADS</code>	<code>_POSIX_THREADS</code>	Supports Threads option
<code>_SC_TIMEOUTS</code>	<code>_POSIX_TIMEOUTS</code>	Supports Timeouts option
<code>_SC_TIMER_MAX</code>	<code>TIMER_MAX</code>	Max number of timer per process supported
<code>_SC_TIMERS</code>	<code>_POSIX_TIMERS</code>	Supports Timers
<code>_SC_TTY_NAME_MAX</code>	<code>TTYNAME_MAX</code>	Max length of tty device name
<code>_SC_TZNAME_MAX</code>	<code>TZNAME_MAX</code>	Max number of bytes supported for name of a time zone
<code>_SC_V6_ILP32_OFF32</code>	<code>_POSIX_V6_ILP32_OFF32</code>	Supports X/Open ILP32 w/32-bit offset build environment
<code>_SC_V6_ILP32_OFFBIG</code>	<code>_POSIX_V6_ILP32_OFFBIG</code>	Supports X/Open ILP32 w/64-bit offset build environment
<code>_SC_V6_LP64_OFF64</code>	<code>_POSIX_V6_LP64_OFF64</code>	Supports X/Open LP64 w/64-bit offset build environment
<code>_SC_V6_LPBIG_OFFBIG</code>	<code>_POSIX_V6_LPBIG_OFFBIG</code>	Same as <code>_SC_V6_LP64_OFF64</code>
<code>_SC_VERSION</code>	<code>_POSIX_VERSION</code>	POSIX.1 version supported
<code>_SC_XBS5_ILP32_OFF32</code>	<code>_XBS_ILP32_OFF32</code>	Indicates support for X/Open ILP32 w/32-bit offset build environment
<code>_SC_XBS5_ILP32_OFFBIG</code>	<code>_XBS5_ILP32_OFFBIG</code>	Indicates support for X/Open ILP32 w/64-bit offset build environment
<code>_SC_XBS5_LP64_OFF64</code>	<code>_XBS5_LP64_OFF64</code>	Indicates support of X/Open LP64, 64-bit offset

sysconf(3C)

<code>_SC_XBS5_LPBIG_OFFBIG</code>	<code>_XBS5_LP64_OFF64</code>	build environment Same as <code>_SC_XBS5_LP64_OFF64</code>
<code>_SC_XOPEN_CRYPT</code>	<code>_XOPEN_CRYPT</code>	Supports X/Open Encryption Feature Group
<code>_SC_XOPEN_ENH_I18N</code>	<code>_XOPEN_ENH_I18N</code>	Supports X/Open Enhanced Internationalization Feature Group
<code>_SC_XOPEN_LEGACY</code>	<code>_XOPEN_LEGACY</code>	Supports X/Open Legacy Feature Group
<code>_SC_XOPEN_REALTIME</code>	<code>_XOPEN_REALTIME</code>	Supports X/Open POSIX Realtime Feature Group
<code>_SC_XOPEN_REALTIME_THREADS</code>	<code>_XOPEN_REALTIME_THREADS</code>	Supports X/Open POSIX Realtime Threads Feature Group
<code>_SC_XOPEN_SHM</code>	<code>_XOPEN_SHM</code>	Supports X/Open Shared Memory Feature Group
<code>_SC_XOPEN_STREAMS</code>	<code>_POSIX_XOPEN_STREAMS</code>	Supports XSI Streams option group
<code>_SC_XOPEN_UNIX</code>	<code>_XOPEN_UNIX</code>	Supports X/Open CAE Specification, August 1994, System Interfaces and Headers, Issue 4, Version 2
<code>_SC_XOPEN_VERSION</code>	<code>_XOPEN_VERSION</code>	Integer value indicates version of X/Open Portability Guide to which implementation conforms
<code>_SC_XOPEN_XCU_VERSION</code>	<code>_XOPEN_XCU_VERSION</code>	Integer value indicates version of XCU specification to which implementation conforms

The following options are not supported and return -1:

<code>_SC_2_PBS</code>	<code>_POSIX2_PBS</code>
<code>_SC_2_PBS_ACCOUNTING</code>	<code>_POSIX2_PBS_ACCOUNTING</code>
<code>_SC_2_PBS_CHECKPOINT</code>	<code>_POSIX2_PBS_CHECKPOINT</code>
<code>_SC_2_PBS_LOCATE</code>	<code>_POSIX2_PBS_LOCATE</code>
<code>_SC_2_PBS_MESSAGE</code>	<code>_POSIX2_PBS_MESSAGE</code>
<code>_SC_2_PBS_TRACK</code>	<code>_POSIX2_PBS_TRACK</code>

<code>_SC_ADVISORY_INFO</code>	<code>_POSIX_ADVISORY_INFO</code>
<code>_SC_CPUTIME</code>	<code>_POSIX_CPUTIME</code>
<code>_SC_SPORADIC_SERVER</code>	<code>_POSIX_SPORADIC_SERVER</code>
<code>_SC_SS_REPL_MAX</code>	<code>_POSIX_SS_REPL_MAX</code>
<code>_SC_THREAD_CPUTIME</code>	<code>_POSIX_THREAD_CPUTIME</code>
<code>_SC_THREAD_SPORADIC_SERVER</code>	<code>_POSIX_THREAD_SPORADIC_SERVER</code>
<code>_SC_TRACE</code>	<code>_POSIX_TRACE</code>
<code>_SC_TRACE_EVENT_FILTER</code>	<code>_POSIX_TRACE_EVENT_FILTER</code>
<code>_SC_TRACE_EVENT_NAME_MAX</code>	<code>_POSIX_TRACE_EVENT_NAME_MAX</code>
<code>_SC_TRACE_INHERIT</code>	<code>_POSIX_TRACE_INHERIT</code>
<code>_SC_TRACE_LOG</code>	<code>_POSIX_TRACE_LOG</code>
<code>_SC_TRACE_NAME_MAX</code>	<code>_POSIX_TRACE_NAME_MAX</code>
<code>_SC_TRACE_SYS_MAX</code>	<code>_POSIX_TRACE_SYS_MAX</code>
<code>_SC_TRACE_USER_EVENT_MAX</code>	<code>_POSIX_TRACE_USER_EVENT_MAX</code>
<code>_SC_TYPED_MEMORY_OBJECTS</code>	<code>_POSIX_TYPED_MEMORY_OBJECTS</code>

RETURN VALUES

Upon successful completion, `sysconf()` returns the current variable value on the system. The value returned will not be more restrictive than the corresponding value described to the application when it was compiled with the implementation's `<limits.h>`, `<unistd.h>` or `<time.h>`. With only a few obvious exceptions such as `_SC_AVPHYS_PAGES` and `_SC_NPROCESSORS_ONLN`, the value will not change during the lifetime of the calling process.

If *name* is an invalid value, `sysconf()` returns `-1` and sets `errno` to indicate the error. If the variable corresponding to *name* is associated with functionality that is not supported by the system, `sysconf()` returns `-1` without changing the value of `errno`.

Calling `sysconf()` with the following returns `-1` without setting `errno`, because no maximum limit can be determined. The system supports at least the minimum values and can support higher values depending upon system resources.

Variable	Minimum supported value
<code>_SC_AIO_MAX</code>	<code>_POSIX_AIO_MAX</code>
<code>_SC_ATEXIT_MAX</code>	32
<code>_SC_THREAD_THREADS_MAX</code>	<code>_POSIX_THREAD_THREADS_MAX</code>
<code>_SC_THREAD_KEYS_MAX</code>	<code>_POSIX_THREAD_KEYS_MAX</code>
<code>_SC_THREAD_DESTRUCTOR_ITERATIONS</code>	<code>_POSIX_THREAD_DESTRUCTOR_ITERATIONS</code>

The following SPARC and x86 platform variables return `EINVAL`:

<code>_SC_COHER_BLKSZ</code>	<code>_SC_DCACHE_ASSOC</code>
<code>_SC_DCACHE_BLKSZ</code>	<code>_SC_DCACHE_LINESZ</code>
<code>_SC_DCACHE_SZ</code>	<code>_SC_DCACHE_TBLKSZ</code>

sysconf(3C)

`_SC_ICACHE_ASSOC` `_SC_ICACHE_BLKSZ`
`_SC_ICACHE_LINESZ` `_SC_ICACHE_SZ`
`_SC_SPLIT_CACHE`

ERRORS The `sysconf()` function will fail if:
`EINVAL` The value of the *name* argument is invalid.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC and x86
Interface Stability	Standard
MT-Level	MT-Safe, Async-Signal-Safe

SEE ALSO `pooladm(1M)`, `zoneadm(1M)`, `fpathconf(2)`, `seteuid(2)`, `setrlimit(2)`, `confstr(3C)`, `attributes(5)`, `standards(5)`

NOTES A call to `setrlimit()` can cause the value of `OPEN_MAX` to change.

Multiplying `sysconf(_SC_PHYS_PAGES)` or `sysconf(_SC_AVPHYS_PAGES)` by `sysconf(_SC_PAGESIZE)` to determine memory amount in bytes can exceed the maximum values representable in a 32-bit signed or unsigned integer.

The value of `CLK_TCK` can be variable and it should not be assumed that `CLK_TCK` is a compile-time constant.

If the caller is in a non-global zone and the pools facility is active, `sysconf(_SC_NPROCESSORS_CONF)` and `sysconf(_SC_NPROCESSORS_ONLN)` return the number of processors in the processor set of the pool to which the zone is bound.

NAME	syslog, openlog, closelog, setlogmask – control system log																
SYNOPSIS	<pre>#include <syslog.h> void openlog(const char *ident, int logopt, int facility); void syslog(int priority, const char *message, .../* arguments */); void closelog(void); int setlogmask(int maskpri);</pre>																
DESCRIPTION	<p>The <code>syslog()</code> function sends a message to <code>syslogd(1M)</code>, which, depending on the configuration of <code>/etc/syslog.conf</code>, logs it in an appropriate system log, writes it to the system console, forwards it to a list of users, or forwards it to <code>syslogd</code> on another host over the network. The logged message includes a message header and a message body. The message header consists of a facility indicator, a severity level indicator, a timestamp, a tag string, and optionally the process ID.</p> <p>The message body is generated from the <i>message</i> and following arguments in the same manner as if these were arguments to <code>printf(3UCB)</code>, except that occurrences of <code>%m</code> in the format string pointed to by the <i>message</i> argument are replaced by the error message string associated with the current value of <code>errno</code>. A trailing <code>NEWLINE</code> character is added if needed.</p> <p>Symbolic constants for use as values of the <i>logopt</i>, <i>facility</i>, <i>priority</i>, and <i>maskpri</i> arguments are defined in the <code><syslog.h></code> header.</p> <p>Values of the <i>priority</i> argument are formed by ORing together a <i>severity level</i> value and an optional <i>facility</i> value. If no facility value is specified, the current default facility value is used.</p> <p>Possible values of severity level include, in decreasing order:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;"><code>LOG_EMERG</code></td> <td>A panic condition. This is normally broadcast to all users.</td> </tr> <tr> <td><code>LOG_ALERT</code></td> <td>A condition that should be corrected immediately, such as a corrupted system database.</td> </tr> <tr> <td><code>LOG_CRIT</code></td> <td>Critical conditions, such as hard device errors.</td> </tr> <tr> <td><code>LOG_ERR</code></td> <td>Errors.</td> </tr> <tr> <td><code>LOG_WARNING</code></td> <td>Warning messages.</td> </tr> <tr> <td><code>LOG_NOTICE</code></td> <td>Conditions that are not error conditions, but that may require special handling.</td> </tr> <tr> <td><code>LOG_INFO</code></td> <td>Informational messages.</td> </tr> <tr> <td><code>LOG_DEBUG</code></td> <td>Messages that contain information normally of use only when debugging a program.</td> </tr> </table>	<code>LOG_EMERG</code>	A panic condition. This is normally broadcast to all users.	<code>LOG_ALERT</code>	A condition that should be corrected immediately, such as a corrupted system database.	<code>LOG_CRIT</code>	Critical conditions, such as hard device errors.	<code>LOG_ERR</code>	Errors.	<code>LOG_WARNING</code>	Warning messages.	<code>LOG_NOTICE</code>	Conditions that are not error conditions, but that may require special handling.	<code>LOG_INFO</code>	Informational messages.	<code>LOG_DEBUG</code>	Messages that contain information normally of use only when debugging a program.
<code>LOG_EMERG</code>	A panic condition. This is normally broadcast to all users.																
<code>LOG_ALERT</code>	A condition that should be corrected immediately, such as a corrupted system database.																
<code>LOG_CRIT</code>	Critical conditions, such as hard device errors.																
<code>LOG_ERR</code>	Errors.																
<code>LOG_WARNING</code>	Warning messages.																
<code>LOG_NOTICE</code>	Conditions that are not error conditions, but that may require special handling.																
<code>LOG_INFO</code>	Informational messages.																
<code>LOG_DEBUG</code>	Messages that contain information normally of use only when debugging a program.																

syslog(3C)

The facility indicates the application or system component generating the message. Possible facility values include:

LOG_KERN	Messages generated by the kernel. These cannot be generated by any user processes.
LOG_USER	Messages generated by random user processes. This is the default facility identifier if none is specified.
LOG_MAIL	The mail system.
LOG_DAEMON	System daemons, such as <code>in.ftpd(1M)</code> .
LOG_AUTH	The authentication / security / authorization system: <code>login(1)</code> , <code>su(1M)</code> , <code>getty(1M)</code> .
LOG_LPR	The line printer spooling system: <code>lpr(1B)</code> , <code>lpc(1B)</code> .
LOG_NEWS	Designated for the USENET network news system.
LOG_UUCP	Designated for the UUCP system; it does not currently use <code>syslog()</code> .
LOG_CRON	The cron/at facility; <code>crontab(1)</code> , <code>at(1)</code> , <code>cron(1M)</code> .
LOG_AUDIT	The audit facility, for example, <code>auditd(1M)</code> .
LOG_LOCAL0	Designated for local use.
LOG_LOCAL1	Designated for local use.
LOG_LOCAL2	Designated for local use.
LOG_LOCAL3	Designated for local use.
LOG_LOCAL4	Designated for local use.
LOG_LOCAL5	Designated for local use.
LOG_LOCAL6	Designated for local use.
LOG_LOCAL7	Designated for local use.

The `openlog()` function sets process attributes that affect subsequent calls to `syslog()`. The *ident* argument is a string that is prepended to every message. The *logopt* argument indicates logging options. Values for *logopt* are constructed by a bitwise-inclusive OR of zero or more of the following:

LOG_PID	Log the process ID with each message. This is useful for identifying specific daemon processes (for daemons that fork).
LOG_CONS	Write messages to the system console if they cannot be sent to <code>syslogd(1M)</code> . This option is safe to use in daemon processes that have no controlling terminal, since <code>syslog()</code> forks before opening the console.

LOG_NDELAY	Open the connection to <code>syslogd(1M)</code> immediately. Normally the open is delayed until the first message is logged. This is useful for programs that need to manage the order in which file descriptors are allocated.
LOG_ODELAY	Delay open until <code>syslog()</code> is called.
LOG_NOWAIT	Do not wait for child processes that have been forked to log messages onto the console. This option should be used by processes that enable notification of child termination using <code>SIGCHLD</code> , since <code>syslog()</code> may otherwise block waiting for a child whose exit status has already been collected.

The *facility* argument encodes a default facility to be assigned to all messages that do not have an explicit facility already encoded. The initial default facility is `LOG_USER`.

The `openlog()` and `syslog()` functions may allocate a file descriptor. It is not necessary to call `openlog()` prior to calling `syslog()`.

The `closelog()` function closes any open file descriptors allocated by previous calls to `openlog()` or `syslog()`.

The `setlogmask()` function sets the log priority mask for the current process to *maskpri* and returns the previous mask. If the *maskpri* argument is 0, the current log mask is not modified. Calls by the current process to `syslog()` with a priority not set in *maskpri* are rejected. The mask for an individual priority *pri* is calculated by the macro `LOG_MASK(pri)`; the mask for all priorities up to and including *toppri* is given by the macro `LOG_UPTO(toppri)`. The default log mask allows all priorities to be logged.

RETURN VALUES The `setlogmask()` function returns the previous log priority mask. The `closelog()`, `openlog()` and `syslog()` functions return no value.

ERRORS No errors are defined.

EXAMPLES **EXAMPLE 1** Example of `LOG_ALERT` message.

This call logs a message at priority `LOG_ALERT`:

```
syslog(LOG_ALERT, "who: internal error 23");
```

The FTP daemon `ftpd` would make this call to `openlog()` to indicate that all messages it logs should have an identifying string of `ftpd`, should be treated by `syslogd(1M)` as other messages from system daemons are, should include the process ID of the process logging the message:

```
openlog("ftpd", LOG_PID, LOG_DAEMON);
```

syslog(3C)

EXAMPLE 1 Example of LOG_ALERT message. *(Continued)*

Then it would make the following call to `setlogmask()` to indicate that messages at priorities from LOG_EMERG through LOG_ERR should be logged, but that no messages at any other priority should be logged:

```
setlogmask(LOG_UPTO(LOG_ERR));
```

Then, to log a message at priority LOG_INFO, it would make the following call to `syslog`:

```
syslog(LOG_INFO, "Connection from host %d", CallingHost);
```

A locally-written utility could use the following call to `syslog()` to log a message at priority LOG_INFO to be treated by `syslogd(1M)` as other messages to the facility LOG_LOCAL2 are:

```
syslog(LOG_INFO|LOG_LOCAL2, "error: %m");
```

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Safe

SEE ALSO

`at(1)`, `crontab(1)`, `logger(1)`, `login(1)`, `lpc(1B)`, `lpr(1B)`, `auditd(1M)`, `cron(1M)`, `getty(1M)`, `in.ftpd(1M)`, `su(1M)`, `syslogd(1M)`, `printf(3UCB)`, `syslog.conf(4)`, `attributes(5)`, `standards(5)`

NAME	system – issue a shell command						
SYNOPSIS	<pre>#include <stdlib.h> int system(const char *string);</pre>						
DESCRIPTION	<p>The <code>system()</code> function causes <i>string</i> to be given to the shell as input, as if <i>string</i> had been typed as a command at a terminal. The invoker waits until the shell has completed, then returns the exit status of the shell in the format specified by <code>waitpid(3C)</code>.</p> <p>If <i>string</i> is a null pointer, <code>system()</code> checks if the shell exists and is executable. If the shell is available, <code>system()</code> returns a non-zero value; otherwise, it returns 0. The standard to which the caller conforms determines which shell is used. See <code>standards(5)</code>.</p>						
RETURN VALUES	The <code>system()</code> function executes <code>vfork(2)</code> to create a child process that in turn invokes one of the <code>exec</code> family of functions (see <code>exec(2)</code>) on the shell to execute <i>string</i> . If <code>vfork()</code> or the <code>exec</code> function fails, <code>system()</code> returns <code>-1</code> and sets <code>errno</code> to indicate the error.						
ERRORS	<p>The <code>system()</code> function fails if:</p> <p>EAGAIN The system-imposed limit on the total number of processes under execution by a single user would be exceeded.</p> <p>EINTR The <code>system()</code> function was interrupted by a signal.</p> <p>ENOMEM The new process requires more memory than is available.</p>						
USAGE	The <code>system()</code> function manipulates the signal handlers for <code>SIGINT</code> , <code>SIGQUIT</code> , and <code>SIGCHLD</code> . For this reason it is not safe to call <code>system()</code> in a multithreaded process. Concurrent calls to <code>system()</code> will interfere destructively with the disposition of these signals, even if they are not manipulated by other threads in the application. See <code>popen(3C)</code> for a replacement for <code>system()</code> that is thread-safe.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>Unsafe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Unsafe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Unsafe						
SEE ALSO	<code>ksh(1)</code> , <code>sh(1)</code> , <code>exec(2)</code> , <code>vfork(2)</code> , <code>popen(3C)</code> , <code>waitpid(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

tcdrain(3C)

NAME	tcdrain – wait for transmission of output								
SYNOPSIS	<pre>#include <termios.h> int tcdrain(int fildes);</pre>								
DESCRIPTION	<p>The <code>tcdrain()</code> function waits until all output written to the object referred to by <i>fildes</i> is transmitted. The <i>fildes</i> argument is an open file descriptor associated with a terminal.</p> <p>Any attempts to use <code>tcdrain()</code> from a process which is a member of a background process group on a <i>fildes</i> associated with its controlling terminal, will cause the process group to be sent a SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process is allowed to perform the operation, and no signal is sent.</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 <code>tcdrain()</code> function will fail if:</p> <table><tr><td>EBADF</td><td>The <i>fildes</i> argument is not a valid file descriptor.</td></tr><tr><td>EINTR</td><td>A signal interrupted <code>tcdrain()</code>.</td></tr><tr><td>ENOTTY</td><td>The file associated with <i>fildes</i> is not a terminal.</td></tr></table> <p>The <code>tcdrain()</code> function may fail if:</p> <table><tr><td>EIO</td><td>The process group of the writing process is orphaned, and the writing process is not ignoring or blocking SIGTTOU.</td></tr></table>	EBADF	The <i>fildes</i> argument is not a valid file descriptor.	EINTR	A signal interrupted <code>tcdrain()</code> .	ENOTTY	The file associated with <i>fildes</i> is not a terminal.	EIO	The process group of the writing process is orphaned, and the writing process is not ignoring or blocking SIGTTOU.
EBADF	The <i>fildes</i> argument is not a valid file descriptor.								
EINTR	A signal interrupted <code>tcdrain()</code> .								
ENOTTY	The file associated with <i>fildes</i> is not a terminal.								
EIO	The process group of the writing process is orphaned, and the writing process is not ignoring or blocking SIGTTOU.								
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:								
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe, and Async-Signal-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe, and Async-Signal-Safe		
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Interface Stability	Standard								
MT-Level	MT-Safe, and Async-Signal-Safe								
SEE ALSO	<code>tcflush(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code> , <code>termio(7I)</code>								

NAME	tcfLOW – suspend or restart the transmission or reception of data								
SYNOPSIS	<pre>#include <termios.h> int tcfLOW(int fildes, int action);</pre>								
DESCRIPTION	<p>The tcfLOW() function suspends transmission or reception of data on the object referred to by <i>fildes</i>, depending on the value of <i>action</i>. The <i>fildes</i> argument is an open file descriptor associated with a terminal.</p> <ul style="list-style-type: none"> ■ If <i>action</i> is TCOFF, output is suspended. ■ If <i>action</i> is TCOON, suspended output is restarted. ■ If <i>action</i> is TCIOFF, the system transmits a STOP character, which is intended to cause the terminal device to stop transmitting data to the system. ■ If <i>action</i> is TCION, the system transmits a START character, which is intended to cause the terminal device to start transmitting data to the system. <p>The default on the opening of a terminal file is that neither its input nor its output are suspended.</p> <p>Attempts to use tcfLOW() from a process which is a member of a background process group on a <i>fildes</i> associated with its controlling terminal, will cause the process group to be sent a SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process is allowed to perform the operation, and no signal is sent.</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 tcfLOW() function will fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EBADF</td> <td>The <i>fildes</i> argument is not a valid file descriptor.</td> </tr> <tr> <td>EINVAL</td> <td>The <i>action</i> argument is not a supported value.</td> </tr> <tr> <td>ENOTTY</td> <td>The file associated with <i>fildes</i> is not a terminal.</td> </tr> </table> <p>The tcfLOW() function may fail if:</p> <table border="0"> <tr> <td style="padding-right: 20px;">EIO</td> <td>The process group of the writing process is orphaned, and the writing process is not ignoring or blocking SIGTTOU.</td> </tr> </table>	EBADF	The <i>fildes</i> argument is not a valid file descriptor.	EINVAL	The <i>action</i> argument is not a supported value.	ENOTTY	The file associated with <i>fildes</i> is not a terminal.	EIO	The process group of the writing process is orphaned, and the writing process is not ignoring or blocking SIGTTOU.
EBADF	The <i>fildes</i> argument is not a valid file descriptor.								
EINVAL	The <i>action</i> argument is not a supported value.								
ENOTTY	The file associated with <i>fildes</i> is not a terminal.								
EIO	The process group of the writing process is orphaned, and the writing process is not ignoring or blocking SIGTTOU.								
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:								
	<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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe, and Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe, and Async-Signal-Safe		
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Interface Stability	Standard								
MT-Level	MT-Safe, and Async-Signal-Safe								
SEE ALSO	tcsendbreak(3C) , attributes(5) , standards(5) , termio(7I)								

tcflush(3C)

NAME	tcflush – flush non-transmitted output data, non-read input data or both								
SYNOPSIS	<pre>#include <termios.h> int tcflush(int fildes, int queue_selector);</pre>								
DESCRIPTION	<p>Upon successful completion, <code>tcflush()</code> discards data written to the object referred to by <i>fildes</i> (an open file descriptor associated with a terminal) but not transmitted, or data received but not read, depending on the value of <i>queue_selector</i>:</p> <ul style="list-style-type: none">■ If <i>queue_selector</i> is <code>TCIFLUSH</code> it flushes data received but not read.■ If <i>queue_selector</i> is <code>TCOFLUSH</code> it flushes data written but not transmitted.■ If <i>queue_selector</i> is <code>TCIOFLUSH</code> it flushes both data received but not read and data written but not transmitted. <p>Attempts to use <code>tcflush()</code> from a process which is a member of a background process group on a <i>fildes</i> associated with its controlling terminal, will cause the process group to be sent a <code>SIGTTOU</code> signal. If the calling process is blocking or ignoring <code>SIGTTOU</code> signals, the process is allowed to perform the operation, and no signal is sent.</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 <code>tcflush()</code> function will fail if:</p> <table><tr><td><code>EBADF</code></td><td>The <i>fildes</i> argument is not a valid file descriptor.</td></tr><tr><td><code>EINVAL</code></td><td>The <i>queue_selector</i> argument is not a supported value.</td></tr><tr><td><code>ENOTTY</code></td><td>The file associated with <i>fildes</i> is not a terminal.</td></tr></table> <p>The <code>tcflush()</code> function may fail if:</p> <table><tr><td><code>EIO</code></td><td>The process group of the writing process is orphaned, and the writing process is not ignoring or blocking <code>SIGTTOU</code>.</td></tr></table>	<code>EBADF</code>	The <i>fildes</i> argument is not a valid file descriptor.	<code>EINVAL</code>	The <i>queue_selector</i> argument is not a supported value.	<code>ENOTTY</code>	The file associated with <i>fildes</i> is not a terminal.	<code>EIO</code>	The process group of the writing process is orphaned, and the writing process is not ignoring or blocking <code>SIGTTOU</code> .
<code>EBADF</code>	The <i>fildes</i> argument is not a valid file descriptor.								
<code>EINVAL</code>	The <i>queue_selector</i> argument is not a supported value.								
<code>ENOTTY</code>	The file associated with <i>fildes</i> is not a terminal.								
<code>EIO</code>	The process group of the writing process is orphaned, and the writing process is not ignoring or blocking <code>SIGTTOU</code> .								
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:								
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe, and Async-Signal-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe, and Async-Signal-Safe		
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
Interface Stability	Standard								
MT-Level	MT-Safe, and Async-Signal-Safe								
SEE ALSO	<code>tcdrain(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code> , <code>termio(7I)</code>								

NAME	tcgetattr – get the parameters associated with the terminal						
SYNOPSIS	<pre>#include <termios.h> int tcgetattr(int <i>fildev</i>, struct termios *<i>termios_p</i>);</pre>						
DESCRIPTION	<p>The <code>tcgetattr()</code> function gets the parameters associated with the terminal referred to by <i>fildev</i> and stores them in the <code>termios</code> structure (see <code>termio(7I)</code>) referenced by <i>termios_p</i>. The <i>fildev</i> argument is an open file descriptor associated with a terminal.</p> <p>The <i>termios_p</i> argument is a pointer to a <code>termios</code> structure.</p> <p>The <code>tcgetattr()</code> operation is allowed from any process.</p> <p>If the terminal device supports different input and output baud rates, the baud rates stored in the <code>termios</code> structure returned by <code>tcgetattr()</code> reflect the actual baud rates, even if they are equal. If differing baud rates are not supported, the rate returned as the output baud rate is the actual baud rate. If the terminal device does not support split baud rates, the input baud rate stored in the <code>termios</code> structure will be 0.</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 <code>tcgetattr()</code> function will fail if:</p> <p><code>EBADF</code> The <i>fildev</i> argument is not a valid file descriptor.</p> <p><code>ENOTTY</code> The file associated with <i>fildev</i> is not a terminal.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe, and Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe, and Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe, and Async-Signal-Safe						
SEE ALSO	<code>tcsetattr(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code> , <code>termio(7I)</code>						

tcgetpgrp(3C)

NAME	tcgetpgrp – get foreground process group ID						
SYNOPSIS	<pre>#include <sys/types.h> #include <unistd.h> pid_t tcgetpgrp (int <i>fildev</i>);</pre>						
DESCRIPTION	<p>The tcgetpgrp () function will return the value of the process group ID of the foreground process group associated with the terminal.</p> <p>If there is no foreground process group, tcgetpgrp () returns a value greater than 1 that does not match the process group ID of any existing process group.</p> <p>The tcgetpgrp () function is allowed from a process that is a member of a background process group; however, the information may be subsequently changed by a process that is a member of a foreground process group.</p>						
RETURN VALUES	Upon successful completion, tcgetpgrp () returns the value of the process group ID of the foreground process associated with the terminal. Otherwise, -1 is returned and errno is set to indicate the error.						
ERRORS	The tcgetpgrp () function will fail if: EBADF The <i>fildev</i> argument is not a valid file descriptor. ENOTTY The calling process does not have a controlling terminal, or the file is not the controlling terminal.						
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>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe, and Async-Signal-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe, and Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe, and Async-Signal-Safe						
SEE ALSO	setpgid(2), setsid(2), tcsetpgrp(3C), attributes(5), standards(5), termio(71)						

- NAME** tcgetsid – get process group ID for session leader for controlling terminal
- SYNOPSIS**

```
#include <termios.h>
pid_t tcgetsid(int fildes);
```
- DESCRIPTION** The `tcgetsid()` function obtains the process group ID of the session for which the terminal specified by *fildes* is the controlling terminal.
- RETURN VALUES** Upon successful completion, `tcgetsid()` returns the process group ID associated with the terminal. Otherwise, a value of `(pid_t)-1` is returned and `errno` is set to indicate the error.
- ERRORS** The `tcgetsid()` function will fail if:
- `EACCES` The *fildes* argument is not associated with a controlling terminal.
 - `EBADF` The *fildes* argument is not a valid file descriptor.
 - `ENOTTY` The file associated with *fildes* is not a terminal.
- ATTRIBUTES** See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

- SEE ALSO** `attributes(5)`, `standards(5)`, `termio(7I)`

tcsendbreak(3C)

NAME	tcsendbreak – send a “break” for a specific duration						
SYNOPSIS	<pre>#include <termios.h> int tcsendbreak(int <i>fildev</i>, int <i>duration</i>);</pre>						
DESCRIPTION	<p>The <i>fildev</i> argument is an open file descriptor associated with a terminal.</p> <p>If the terminal is using asynchronous serial data transmission, <code>tcsendbreak()</code> will cause transmission of a continuous stream of zero-valued bits for a specific duration. If <i>duration</i> is 0, it will cause transmission of zero-valued bits for at least 0.25 seconds, and not more than 0.5 seconds. If <i>duration</i> is not 0, it behaves in a way similar to <code>tcdrain(3C)</code>.</p> <p>If the terminal is not using asynchronous serial data transmission, it sends data to generate a break condition or returns without taking any action.</p> <p>Attempts to use <code>tcsendbreak()</code> from a process which is a member of a background process group on a <i>fildev</i> associated with its controlling terminal will cause the process group to be sent a SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process is allowed to perform the operation, and no signal is sent.</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 <code>tcsendbreak()</code> function will fail if:</p> <table><tr><td>EBADF</td><td>The <i>fildev</i> argument is not a valid file descriptor.</td></tr><tr><td>ENOTTY</td><td>The file associated with <i>fildev</i> is not a terminal.</td></tr></table> <p>The <code>tcsendbreak()</code> function may fail if:</p> <table><tr><td>EIO</td><td>The process group of the writing process is orphaned, and the writing process is not ignoring or blocking SIGTTOU.</td></tr></table>	EBADF	The <i>fildev</i> argument is not a valid file descriptor.	ENOTTY	The file associated with <i>fildev</i> is not a terminal.	EIO	The process group of the writing process is orphaned, and the writing process is not ignoring or blocking SIGTTOU.
EBADF	The <i>fildev</i> argument is not a valid file descriptor.						
ENOTTY	The file associated with <i>fildev</i> is not a terminal.						
EIO	The process group of the writing process is orphaned, and the writing process is not ignoring or blocking SIGTTOU.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe, and Async-Signal-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe, and Async-Signal-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe, and Async-Signal-Safe						
SEE ALSO	<code>tcdrain(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code> , <code>termio(7I)</code>						

NAME	tcsetattr – set the parameters associated with the terminal
SYNOPSIS	<pre>#include <termios.h> int tcsetattr(int fildes, int optional_actions, const struct termios *termios_p);</pre>
DESCRIPTION	<p>The <code>tcsetattr()</code> function sets the parameters associated with the terminal referred to by the open file descriptor <i>fildes</i> (an open file descriptor associated with a terminal) from the <code>termios</code> structure (see <code>termio(7I)</code>) referenced by <i>termios_p</i> as follows:</p> <ul style="list-style-type: none"> ■ If <i>optional_actions</i> is <code>TCSANOW</code>, the change will occur immediately. ■ If <i>optional_actions</i> is <code>TCSADRAIN</code>, the change will occur after all output written to <i>fildes</i> is transmitted. This function should be used when changing parameters that affect output. ■ If <i>optional_actions</i> is <code>TCSAFLUSH</code>, the change will occur after all output written to <i>fildes</i> is transmitted, and all input so far received but not read will be discarded before the change is made. <p>If the output baud rate stored in the <code>termios</code> structure pointed to by <i>termios_p</i> is the zero baud rate, <code>B0</code>, the modem control lines will no longer be asserted. Normally, this will disconnect the line.</p> <p>If the input baud rate stored in the <code>termios</code> structure pointed to by <i>termios_p</i> is 0, the input baud rate given to the hardware will be the same as the output baud rate stored in the <code>termios</code> structure.</p> <p>The <code>tcsetattr()</code> function will return successfully if it was able to perform any of the requested actions, even if some of the requested actions could not be performed. It will set all the attributes that implementation supports as requested and leave all the attributes not supported by the implementation unchanged. If no part of the request can be honoured, it will return <code>-1</code> and set <code>errno</code> to <code>EINVAL</code>. If the input and output baud rates differ and are a combination that is not supported, neither baud rate is changed. A subsequent call to <code>tcgetattr(3C)</code> will return the actual state of the terminal device (reflecting both the changes made and not made in the previous <code>tcsetattr()</code> call). The <code>tcsetattr()</code> function will not change the values in the <code>termios</code> structure whether or not it actually accepts them.</p> <p>The effect of <code>tcsetattr()</code> is undefined if the value of the <code>termios</code> structure pointed to by <i>termios_p</i> was not derived from the result of a call to <code>tcgetattr(3C)</code> on <i>fildes</i>; an application should modify only fields and flags defined by this document between the call to <code>tcgetattr(3C)</code> and <code>tcsetattr()</code>, leaving all other fields and flags unmodified.</p> <p>No actions defined by this document, other than a call to <code>tcsetattr()</code> or a close of the last file descriptor in the system associated with this terminal device, will cause any of the terminal attributes defined by this document to change.</p>

tcsetattr(3C)

Attempts to use `tcsetattr()` from a process which is a member of a background process group on a *fildev* associated with its controlling terminal, will cause the process group to be sent a SIGTTOU signal. If the calling process is blocking or ignoring SIGTTOU signals, the process is allowed to perform the operation, and no signal is sent.

USAGE If trying to change baud rates, applications should call `tcsetattr()` then call `tcgetattr(3C)` in order to determine what baud rates were actually selected.

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

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

- `EBADF` The *fildev* argument is not a valid file descriptor.
- `EINTR` A signal interrupted `tcsetattr()`.
- `EINVAL` The *optional_actions* argument is not a supported value, or an attempt was made to change an attribute represented in the `termios` structure to an unsupported value.
- `ENOTTY` The file associated with *fildev* is not a terminal.

The `tcsetattr()` function may fail if:

- `EIO` The process group of the writing process is orphaned, and the writing process is not ignoring or blocking SIGTTOU.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe, and Async-Signal-Safe

SEE ALSO `cfgetispeed(3C)`, `tcgetattr(3C)`, `attributes(5)`, `standards(5)`, `termio(7I)`

NAME	tcsetpgrp – set foreground process group ID										
SYNOPSIS	<pre>#include <sys/types.h> #include <unistd.h> int tcsetpgrp(int <i>fildev</i>, pid_t <i>pgid_id</i>);</pre>										
DESCRIPTION	<p>If the process has a controlling terminal, <code>tcsetpgrp()</code> will set the foreground process group ID associated with the terminal to <code>pgid_id</code>. The file associated with <code>fildev</code> must be the controlling terminal of the calling process and the controlling terminal must be currently associated with the session of the calling process. The value of <code>pgid_id</code> must match a process group ID of a process in the same session as the calling process.</p>										
RETURN VALUES	<p>Upon successful completion, 0 is returned. Otherwise, -1 is returned and <code>errno</code> is set to indicate the error.</p>										
ERRORS	<p>The <code>tcsetpgrp()</code> function will fail if:</p> <table> <tr> <td>EBADF</td> <td>The <code>fildev</code> argument is not a valid file descriptor.</td> </tr> <tr> <td>EINVAL</td> <td>This implementation does not support the value in the <code>pgid_id</code> argument.</td> </tr> <tr> <td>ENOTTY</td> <td>The calling process does not have a controlling terminal, or the file is not the controlling terminal, or the controlling terminal is no longer associated with the session of the calling process.</td> </tr> <tr> <td>EIO</td> <td>The process is not ignoring or holding SIGTTOU and is a member of an orphaned process group.</td> </tr> <tr> <td>EPERM</td> <td>The value of <code>pgid_id</code> does not match the process group ID of a process in the same session as the calling process.</td> </tr> </table>	EBADF	The <code>fildev</code> argument is not a valid file descriptor.	EINVAL	This implementation does not support the value in the <code>pgid_id</code> argument.	ENOTTY	The calling process does not have a controlling terminal, or the file is not the controlling terminal, or the controlling terminal is no longer associated with the session of the calling process.	EIO	The process is not ignoring or holding SIGTTOU and is a member of an orphaned process group.	EPERM	The value of <code>pgid_id</code> does not match the process group ID of a process in the same session as the calling process.
EBADF	The <code>fildev</code> argument is not a valid file descriptor.										
EINVAL	This implementation does not support the value in the <code>pgid_id</code> argument.										
ENOTTY	The calling process does not have a controlling terminal, or the file is not the controlling terminal, or the controlling terminal is no longer associated with the session of the calling process.										
EIO	The process is not ignoring or holding SIGTTOU and is a member of an orphaned process group.										
EPERM	The value of <code>pgid_id</code> does not match the process group ID of a process in the same session as the calling process.										
ATTRIBUTES	<p>See <code>attributes(5)</code> 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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe, and Async-Signal-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe, and Async-Signal-Safe				
ATTRIBUTE TYPE	ATTRIBUTE VALUE										
Interface Stability	Standard										
MT-Level	MT-Safe, and Async-Signal-Safe										
SEE ALSO	<code>tcgetpgrp(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code> , <code>termio(7I)</code>										

td_init(3C_DB)

NAME | td_init – performs initialization for libc_db library of interfaces

SYNOPSIS |

```
cc [ flag... ] file... -lc_db [ library... ]
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_init();
```

DESCRIPTION | The td_init() function is the global initialization function for the libc_db() library of interfaces. It must be called exactly once by any process using the libc_db() library before any other libc_db() function can be called.

RETURN VALUES | TD_OK The libc_db() library of interfaces successfully initialized.
TD_ERR Initialization failed.

ATTRIBUTES | See attributes(5) for description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

SEE ALSO | libc_db(3LIB), attributes(5)

NAME td_log – placeholder for future logging functionality

SYNOPSIS

```
cc [ flag... ] file... -lc_db [ library... ]
#include <proc_service.h>
#include <thread_db.h>

void td_log(void);
```

DESCRIPTION This function presently does nothing. It is merely a placeholder for future logging functionality in libc_db(3LIB).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

SEE ALSO libc_db(3LIB), attributes(5), threads(5)

td_sync_get_info(3C_DB)

NAME	td_sync_get_info, td_ta_sync_tracking_enable, td_sync_get_stats, td_sync_setstate, td_sync_waiters – operations on a synchronization object in libc_db		
SYNOPSIS	<pre>cc [flag...] file... -lc_db [library...] #include <proc_service.h> #include <thread_db.h> td_err_e td_sync_get_info(const td_synchandle_t *sh_p, td_syncinfo_t *si_p); td_err_e td_ta_sync_tracking_enable(const td_thragent_t *ta_p, int on_off); td_err_e td_sync_get_stats(const td_synchandle_t *sh_p, td_syncstats_t *ss_p); td_err_e td_sync_setstate(const td_synchandle_t *sh_p); typedef int td_thr_iter_f(const td_thrhandle_t *th_p, void *cb_data_p); td_err_e td_sync_waiters(const td_synchandle_t *sh_p, td_thr_iter_f *cb, void *cb_data_p);</pre>		
DESCRIPTION	<p>Synchronization objects include mutexes, condition variables, semaphores, and reader-writer locks. In the same way that thread operations use a thread handle of type <code>td_thrhandle_t</code>, operations on synchronization objects use a synchronization object handle of type <code>td_synchandle_t</code>.</p> <p>The controlling process obtains synchronization object handles either by calling the function <code>td_ta_sync_iter()</code> to obtain handles for all synchronization objects of the target process that are known to the <code>libc_db</code> library of interfaces, or by mapping the address of a synchronization object in the address space of the target process to a handle by calling <code>td_ta_map_addr2sync(3C_DB)</code>.</p> <p>Not all synchronization objects that a process uses can be known to the <code>libc_db</code> library and returned by <code>td_ta_sync_iter(3C_DB)</code>. A synchronization object is known to <code>libc_db</code> only if it has been the target of a synchronization primitive in the process (such as <code>mutex_lock()</code>, described on the <code>mutex_init(3C)</code> manual page) after <code>td_ta_new(3C_DB)</code> has been called to attach to the process and <code>td_ta_sync_tracking_enable()</code> has been called to enable synchronization object tracking.</p> <p>The <code>td_ta_sync_tracking_enable()</code> function turns synchronization object tracking on or off for the process identified by <code>ta_p</code>, depending on whether <code>on_off</code> is 0 (off) or non-zero (on).</p> <p>The <code>td_sync_get_info()</code> function fills in the <code>td_syncinfo_t</code> structure <code>*si_p</code> with values for the synchronization object identified by <code>sh_p</code>. The <code>td_syncinfo_t</code> structure contains the following fields:</p> <table><tr><td><code>td_thragent_t *si_ta_p</code></td><td>The internal process handle identifying the target process through which this</td></tr></table>	<code>td_thragent_t *si_ta_p</code>	The internal process handle identifying the target process through which this
<code>td_thragent_t *si_ta_p</code>	The internal process handle identifying the target process through which this		

td_sync_get_info(3C_DB)

	synchronization object handle was obtained. Synchronization objects may be process-private or process-shared. In the latter case, the same synchronization object may have multiple handles, one for each target process's "view" of the synchronization object.
<code>psaddr_t si_sv_addr</code>	The address of the synchronization object in this target process's address space.
<code>td_sync_type_e si_type</code>	The type of the synchronization variable: mutex, condition variable, semaphore, or readers-writer lock.
<code>int si_shared_type</code>	If the <code>USYNC_PROCESS</code> or <code>USYNC_PROCESS_ROBUST</code> bit is set in <code>si_shared_type</code> , this synchronization object is process-shared, otherwise it is process-private.
<code>td_sync_flags_t si_flags</code>	Flags dependent on the type of the synchronization object.
<code>int si_state.sema_count</code>	Semaphores only. The current value of the semaphore
<code>int si_state.nreaders</code>	Readers-writer locks only. The number of readers currently holding the lock, or -1, if a writer is currently holding the lock.
<code>int si_state.mutex_locked</code>	For mutexes only. Non-zero if and only if the mutex is currently locked.
<code>int si_size</code>	The size of the synchronization object.
<code>uint8_t si_has_waiters</code>	Non-zero if and only if at least one thread is blocked on this synchronization object.
<code>uint8_t si_is_wlocked</code>	For reader-writer locks only. The value is non-zero if and only if this lock is held by a writer.
<code>uint8_t si_rcount</code>	<code>PTHREAD_MUTEX_RECURSIVE</code> mutexes only. If the mutex is held, the recursion count.
<code>uint8_t si_priocelling</code>	<code>PTHREAD_PRIO_PROTECT</code> protocol mutexes only. The priority ceiling.

td_sync_get_info(3C_DB)

<code>td_thrhandle_t si_owner</code>	Mutexes and readers-writer locks only. This is the thread holding the mutex, or the write lock, if this is a reader-writer lock. The value is NULL if no one holds the mutex or write-lock.
<code>pid_t si_ownerpid</code>	Mutexes only. For a locked process-shared mutex, this is the process-ID of the process containing the owning thread.

The `td_sync_get_stats()` function fills in the `td_syncstats_t` structure `*ss_p` with values for the synchronization object identified by `sh_p`. The `td_syncstats_t` structure contains an embedded `td_syncinfo_t` structure that is filled in as described above for `td_sync_get_info()`. In addition, usage statistics gathered since `td_ta_sync_tracking_enable()` was called to enable synchronization object tracking are returned in the `ss_un.mutex`, `ss_un.cond`, `ss_un.rwlock`, or `ss_un.sema` members of the `td_syncstats_t` structure, depending on the type of the synchronization object.

The `td_sync_setstate` function modifies the state of synchronization object `si_p`, depending on the synchronization object type. For mutexes, `td_sync_setstate` is unlocked if the value is 0. Otherwise it is locked. For semaphores, the semaphore's count is set to the value. For reader-writer locks, the reader count set to the value if value is >0. The count is set to write-locked if value is -1. It is set to unlocked if the value is 0. Setting the state of a synchronization object from a `libc_db` interface may cause the synchronization object's semantics to be violated from the point of view of the threads in the target process. For example, if a thread holds a mutex, and `td_sync_setstate` is used to set the mutex to unlocked, then a different thread will also be able to subsequently acquire the same mutex.

The `td_sync_waiters` function iterates over the set of thread handles of threads blocked on `sh_p`. The callback function `cb` is called once for each such thread handle, and is passed the thread handle and `cb_data_p`. If the callback function returns a non-zero value, iteration is terminated early. See `td_ta_thr_iter(3C_DB)`.

RETURN VALUES

<code>TD_OK</code>	The call returned successfully.
<code>TD_BADTH</code>	An invalid thread handle was passed in.
<code>TD_DBERR</code>	A call to one of the imported interface routines failed.
<code>TD_ERR</code>	A <code>libc_db</code> -internal error occurred.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

`td_sync_get_info(3C_DB)`

SEE ALSO `libc_db(3LIB)`, `mutex_init(3C)`, `td_ta_map_addr2sync(3C_DB)`,
`td_ta_sync_iter(3C_DB)`, `td_ta_thr_iter(3C_DB)`, `attributes(5)`

td_ta_enable_stats(3C_DB)

NAME	td_ta_enable_stats, td_ta_reset_stats, td_ta_get_stats – collect target process statistics for libc_db
SYNOPSIS	<pre>cc [flag...] file... -lc_db [library...] #include <proc_service.h> #include <thread_db.h> td_err_e td_ta_enable_stats(const td_thragent_t *ta_p, int on_off); td_err_e_stats td_ta_reset(const td_thragent_t *ta_p); td_err_e td_ta_get_stats(const td_thragent_t *ta_p, td_ta_stats_t *tstats);</pre>
DESCRIPTION	<p>The controlling process can request the collection of certain statistics about a target process. Statistics gathering is disabled by default. Each target process has a <code>td_ta_stats_t</code> structure that contains current values when statistic gathering is enabled.</p> <p>The <code>td_ta_enable_stats()</code> function turns statistics gathering on or off for the process identified by <code>ta_p</code>, depending on whether or not <code>on_off</code> is non-zero. When statistics gathering is turned on, all statistics are implicitly reset as though <code>td_ta_reset_stats()</code> had been called. Statistics are not reset when statistics gathering is turned off. Except for <code>nthreads</code> and <code>r_concurrency</code>, the values do not change further, but they remain available for inspection by way of <code>td_ta_get_stats()</code>.</p> <p>The <code>td_ta_reset_stats()</code> function resets all counters in the <code>td_ta_stats_t</code> structure to zero for the target process.</p> <p>The <code>td_ta_get_stats()</code> function returns the structure for the process in <code>tstats</code>.</p> <p>The <code>td_ta_stats_t</code> structure is defined in <code><thread_db.h></code> and contains the following members:</p> <pre>typedef struct { int nthreads; /* total number of threads in use */ int r_concurrency; /* requested concurrency level */ int nrunnable_num; /* numerator of avg runnable threads */ int nrunnable_den; /* denominator of avg runnable threads */ int a_concurrency_num; /* numerator, avg achieved concurrency */ int a_concurrency_den; /* denominator, avg achieved concurrency */ int nlwps_num; /* numerator, avg number of LWPs in use */ int nlwps_den; /* denominator, avg number of LWPs in use */ int nidle_num; /* numerator, avg number of idling LWPs */ int nidle_den; /* denominator, avg number of idling LWPs */ } td_ta_stats_t;</pre> <p>The <code>nthreads</code> member is the number of threads that are currently part of the target process. The <code>r_concurrency</code> member is the current requested concurrency level, such as would be returned by <code>thr_setconcurrency(3C)</code>. The remaining members are averages over time, each expressed as a fraction with an integral numerator and denominator. The <code>nrunnable_num</code> and <code>nrunnable_den</code> members represent the average number of runnable threads. The <code>a_concurrency_num</code> and</p>

td_ta_enable_stats(3C_DB)

a_concurrency_den members represent the average achieved concurrency, the number of actually running threads. The a_concurrency_num and a_concurrency_den members are less than or equal to nrunnable_num and nrunnable_den, respectively. The nlwps_num and nlwps_den members represent the average number of lightweight processes (LWPs) participating in this process. They must be greater than or equal to a_concurrency_num and a_concurrency_den, respectively, since every running thread is assigned to an LWP, but there can at times be additional idling LWPs with no thread assigned to them. The nidle_num and nidle_den members represent the average number of idle LWPs.

RETURN VALUES

TD_OK The call completed successfully.
TD_BADTA An invalid internal process handle was passed in.
TD_DBERR A call to one of the imported interface routines failed.
TD_ERR Something else went wrong.

ATTRIBUTES

See attributes(5) for description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT Level	Safe

SEE ALSO

libc_db(3LIB), thr_getconcurrency(3C), attributes(5)

td_ta_event_addr(3C_DB)

NAME	td_ta_event_addr, td_thr_event_enable, td_ta_set_event, td_thr_set_event, td_ta_clear_event, td_thr_clear_event, td_ta_event_getmsg, td_thr_event_getmsg, td_event_emptyset, td_event_fillset, td_event_addset, td_event_delset, td_eventismember, td_eventisempty – thread events in libc_db
SYNOPSIS	<pre>cc [flag...] file... -lc_db [library...] #include <proc_service.h> #include <thread_db.h> td_err_e td_ta_event_addr(const td_thragent_t *ta_p, u_long event, td_notify_t *notify_p); td_err_e td_thr_event_enable(const td_thrhandle_t *th_p, int on_off); td_err_e td_thr_set_event(const td_thrhandle_t *th_p, td_thr_events_t *events); td_err_e td_ta_set_event(const td_thragent_t *ta_p, td_thr_events_t *events); td_err_e td_thr_clear_event(const td_thrhandle_t *th_p, td_thr_events_t *events); td_err_e td_ta_clear_event(const td_thragent_t *ta_p, td_thr_events_t *events); td_err_e td_thr_event_getmsg(const td_thrhandle_t *th_p, td_event_msg_t *msg); td_err_e td_ta_event_getmsg(const td_thragent_t *ta_p, td_event_msg_t *msg); void td_event_emptyset(td_thr_events_t *); void td_event_fillset(td_thr_events_t *); void td_event_addset(td_thr_events_t *, td_thr_events_e n); void td_event_delset(td_thr_events_t *, td_thr_events_e n); void td_eventismember(td_thr_events_t *, td_thr_events_e n); void td_eventisempty(td_thr_events_t*);</pre>
DESCRIPTION	<p>These functions comprise the thread event facility for libc_db(3LIB). This facility allows the controlling process to be notified when certain thread-related events occur in a target process and to retrieve information associated with these events. An event consists of an event type, and optionally, some associated event data, depending on the event type. See the section titled "Event Set Manipulation Macros" that follows.</p> <p>The event type and the associated event data, if any, constitute an "event message." "Reporting an event" means delivering an event message to the controlling process by way of libc_db.</p>

Several flags can control event reporting, both a per-thread and per event basis. Event reporting may further be enabled or disabled for a thread. There is not only a per-thread event mask that specifies which event types should be reported for that thread, but there is also a global event mask that applies to all threads.

An event is reported, if and only if, the executing thread has event reporting enabled, and either the event type is enabled in the executing thread's event mask, or the event type is enabled in the global event mask.

Each thread has associated with it an event buffer in which it stores the most recent event message it has generated, the type of the most recent event that it reported, and, depending on the event type, some additional information related to that event. See the section titled "Event Set Manipulation Macros" for a description of the `td_thr_events_e` and `td_event_msg_t` types and a list of the event types and the values reported with them. The thread handle, type `td_thrhandle_t`, the event type, and the possible value, together constitute an event message. Each thread's event buffer holds at most one event message.

Each event type has an event reporting address associated with it. A thread reports an event by writing the event message into the thread's event buffer and having control reach the event reporting address for that event type.

Typically, the controlling process sets a breakpoint at the event reporting address for one or more event types. When the breakpoint is hit, the controlling process knows that an event of the corresponding type has occurred.

The event types, and the additional information, if any, reported with each event, are:

<code>TD_READY</code>	The thread became ready to execute.
<code>TD_SLEEP</code>	The thread has blocked on a synchronization object.
<code>TD_SWITCHTO</code>	A runnable thread is being assigned to LWP.
<code>TD_SWITCHFROM</code>	A running thread is being removed from its LWP.
<code>TD_LOCK_TRY</code>	A thread is trying to get an unavailable lock.
<code>TD_CATCHSIG</code>	A signal was posted to a thread.
<code>TD_IDLE</code>	An LWP is becoming idle.
<code>TD_CREATE</code>	A thread is being created.
<code>TD_DEATH</code>	A thread has terminated.
<code>TD_PREEMPT</code>	A thread is being preempted.
<code>TD_PRI_INHERIT</code>	A thread is inheriting an elevated priority from another thread.
<code>TD_REAP</code>	A thread is being reaped.
<code>TD_CONCURRENCY</code>	The number of LWPs is changing.
<code>TD_TIMEOUT</code>	A condition-variable timed wait expired.

td_ta_event_addr(3C_DB)

The `td_ta_event_addr()` function returns in **notify_p* the event reporting address associated with event type *event*. The controlling process may then set a breakpoint at that address. If a thread hits that breakpoint, it reports an event of type *event*.

The `td_thr_event_enable()` function enables or disables event reporting for thread *th_p*. If a thread has event reporting disabled, it will not report any events. Threads are started with event reporting disabled. Event reporting is enabled if *on_off* is non-zero; otherwise, it is disabled. To determine whether or not event reporting is enabled on a thread, call `td_thr_getinfo()` for the thread and examine the `ti_traceme` member of the `td_thrinfo_t` structure it returns.

The `td_thr_set_event()` and `td_thr_clear_event()` functions set and clear, respectively, a set of event types in the event mask associated with the thread *th_p*. To inspect a thread's event mask, call `td_thr_getinfo()` for the thread and examine the `ti_events` member of the `td_thrinfo_t` structure it returns.

The `td_ta_set_event()` and `td_ta_clear_event()` functions identical to `td_thr_set_event()` and `td_thr_clear_event()`, respectively, except that the target process's global event mask is modified. There is no provision for inspecting the value of a target process's global event mask.

The `td_thr_event_getmsg()` function returns in **msg* the event message associated with thread **th_p*. Reading a thread's event message consumes the message, emptying the thread's event buffer. As noted above, each thread's event buffer holds at most one event message; if a thread reports a second event before the first event message has been read, the second event message overwrites the first.

The `td_ta_event_getmsg()` function is identical to `td_thr_event_getmsg()`, except that it is passed a process handle rather than a thread handle. It selects some thread that has an event message buffered and returns that thread's message. The thread selected is undefined, except that as long as at least one thread has an event message buffered, it returns an event message from some such thread.

Event Set Manipulation Macros

Several macros are provided for manipulating event sets of type `td_thr_events_t`:

<code>td_event_emptyset</code>	Sets its argument to the NULL event set.
<code>td_event_fillset</code>	Sets its argument to the set of all events.
<code>td_event_addset</code>	Adds a specific event type to an event set.
<code>td_event_delset</code>	Deletes a specific event type from an event set.
<code>td_eventismember</code>	Tests whether a specific event type is a member of an event set.
<code>td_eventisempty</code>	Tests whether an event set is the NULL set.

RETURN VALUES

The following values may be returned for all thread event routines:

`TD_OK` The call returned successfully.

td_ta_event_addr(3C_DB)

TD_BADTH An invalid thread handle was passed in.
TD_BADTA An invalid internal process handle was passed.
TD_BADPH There is a NULL external process handle associated with this internal process handle.
TD_DBERR A call to one of the imported interface routines failed.
TD_NOMSG No event message was available to return to `td_thr_event_getmsg()` or `td_ta_event_getmsg()`.
TD_ERR Some other parameter error occurred, or a `libc_db()` internal error occurred.

The following value can be returned for `td_thr_event_enable()`, `td_thr_set_event()`, and `td_thr_clear_event()` only:

TD_NOCAPAB Because the agent thread in the target process has not completed initialization, this operation cannot be performed. The operation can be performed after the target process has been allowed to make some forward progress. See `libc_db(3LIB)`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `libc_db(3LIB)`, `attributes(5)`

td_ta_get_nthreads(3C_DB)

NAME | td_ta_get_nthreads – gets the total number of threads in a process for libc_db

SYNOPSIS |

```
cc [ flag... ] file... -lc_db [ library... ]
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_ta_get_nthreads(const td_thragent_t *ta_p, int
    *nthread_p);
```

DESCRIPTION | The `td_ta_get_nthreads()` function returns the total number of threads in process `ta_p`, including any system threads. System threads are those created by `libc` or `libc_db` on its own behalf. The number of threads is written into `*nthread_p`.

RETURN VALUES |

TD_OK	The call completed successfully.
TD_BADTA	An invalid internal process handle was passed in.
TD_BADPH	There is a NULL external process handle associated with this internal process handle.
TD_DBERR	A call to one of the imported interface routines failed.
TD_ERR	The <code>nthread_p</code> argument was NULL, or a <code>libc_db</code> internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO | `libc_db(3LIB)`, `attributes(5)`, `threads(5)`

NAME td_ta_map_addr2sync – get a synchronization object handle from a synchronization object's address

SYNOPSIS

```
cc [ flag... ] file... -lc_db [ library... ]
#include <proc_service.h>
#include <thread_db.h>

td_ta_map_addr2sync(const td_thragent_t *ta_p, psaddr_t addr,
                    td_synchandle_t *sh_p);
```

DESCRIPTION The `td_ta_map_addr2sync()` function produces the synchronization object handle of type `td_synchandle_t` that corresponds to the address of the synchronization object (mutex, semaphore, condition variable, or reader/writer lock). Some effort is made to validate *addr* and verify that it does indeed point at a synchronization object. The handle is returned in **sh_p*.

RETURN VALUES

TD_OK	The call completed successfully.
TD_BADTA	An invalid internal process handle was passed.
TD_BADPH	There is a NULL external process handle associated with this internal process handle.
TD_BADSH	The <i>sh_p</i> argument is NULL or <i>addr</i> does not appear to point to a valid synchronization object.
TD_DBERR	A call to one of the imported interface routines failed.
TD_ERR	<i>addr</i> is NULL, or a <code>libc_db</code> internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `libc_db(3LIB)`, `attributes(5)`

td_ta_map_id2thr(3C_DB)

NAME | td_ta_map_id2thr, td_ta_map_lwp2thr – convert a thread ID or LWP ID to a thread handle

SYNOPSIS |

```
cc [ flag... ] file... -lc_db [ library... ]
#include <proc_service.h>
#include <thread_db.h>

td_ta_map_id2thr(const td_thragent_t *ta_p, thread_t
tid,td_thrhandle_t *th_p);

td_ta_map_lwp2thr(const td_thragent_t *ta_p, lwpid_t
lwpid,td_thrhandle_t *th_p);
```

DESCRIPTION | The td_ta_map_id2thr() function produces the td_thrhandle_t thread handle that corresponds to a particular thread ID, as returned by thr_create(3C) or thr_self(3C). The thread handle is returned in *th_p.

The td_ta_map_lwp2thr() function produces the td_thrhandle_t thread handle for the thread that is currently executing on the light weight process (LWP) and has an ID of lwpid.

RETURN VALUES | TD_OK The call completed successfully.

TD_BADTA An invalid internal process handle was passed in.

TD_BADPH There is a NULL external process handle associated with this internal process handle.

TD_DBERR A call to one of the imported interface routines failed.

TD_NOTHR Either there is no thread with the given thread ID (td_ta_map_id2thr) or no thread is currently executing on the given LWP (td_ta_map_lwp2thr).

TD_ERR The call did not complete successfully.

ATTRIBUTES | See attributes(5) for description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO | libc_db(3LIB), thr_create(3C), thr_self(3C), attributes(5)

NAME	td_ta_new, td_ta_delete, td_ta_get_ph – allocate and deallocate process handles for libc_db												
SYNOPSIS	<pre>cc [flag...] file... -lc_db [library...] #include <proc_service.h> #include <thread_db.h> td_err_e td_ta_new(const struct ps_prochandle *ph_p, td_thragent_t **ta_pp); td_err_e td_ta_delete(const td_thragent_t *ta_p); td_err_e td_ta_get_ph(const td_thragent_t *ta_p, struct ps_prochandle **ph_pp);</pre>												
DESCRIPTION	<p>The <code>td_ta_new()</code> function registers a target process with <code>libc_db(3LIB)</code> and allocates an internal process handle of type <code>td_thragent_t</code> for this target process. Subsequent calls to <code>libc_db</code> can use this handle to refer to this target process.</p> <p>There are actually two process handles, an internal process handle assigned by <code>libc_db</code> and an external process handle assigned by the <code>libc_db</code> client. There is a one-to-one correspondence between the two handles. When the client calls a <code>libc_db</code> function, it uses the internal process handle. When <code>libc_db</code> calls one of the client-provided routines listed in <code>proc_service(3PROC)</code>, it uses the external process handle.</p> <p>The <code>ph</code> argument is the external process handle that <code>libc_db</code> should use to identify this target process to the controlling process when it calls routines in the imported interface.</p> <p>If this call is successful, the value of the newly allocated <code>td_thragent_t</code> handle is returned in <code>*ta_pp</code>. The <code>td_ta_delete()</code> function deregisters a target process with <code>libc_db</code>, which deallocates its internal process handle and frees any other resources <code>libc_db</code> has acquired with respect to the target process. The <code>ta_p</code> argument specifies the target process to be deregistered.</p> <p>The <code>td_ta_get_ph()</code> function returns in <code>*ph_pp</code> the external process handle that corresponds to the internal process handle <code>ta_p</code>. This is useful for checking internal consistency.</p>												
RETURN VALUES	<table border="0"> <tr> <td style="padding-right: 20px;">TD_OK</td> <td>The call completed successfully.</td> </tr> <tr> <td>TD_BADPH</td> <td>A NULL external process handle was passed to <code>td_ta_new()</code>.</td> </tr> <tr> <td>TD_ERR</td> <td>The <code>ta_pp</code> argument is NULL or an internal error occurred.</td> </tr> <tr> <td>TD_DBERR</td> <td>A call to one of the imported interface routines failed.</td> </tr> <tr> <td>TD_MALLOC</td> <td>Memory allocation failure.</td> </tr> <tr> <td>TD_NOLIBTHREAD</td> <td>The target process does not appear to be multithreaded.</td> </tr> </table>	TD_OK	The call completed successfully.	TD_BADPH	A NULL external process handle was passed to <code>td_ta_new()</code> .	TD_ERR	The <code>ta_pp</code> argument is NULL or an internal error occurred.	TD_DBERR	A call to one of the imported interface routines failed.	TD_MALLOC	Memory allocation failure.	TD_NOLIBTHREAD	The target process does not appear to be multithreaded.
TD_OK	The call completed successfully.												
TD_BADPH	A NULL external process handle was passed to <code>td_ta_new()</code> .												
TD_ERR	The <code>ta_pp</code> argument is NULL or an internal error occurred.												
TD_DBERR	A call to one of the imported interface routines failed.												
TD_MALLOC	Memory allocation failure.												
TD_NOLIBTHREAD	The target process does not appear to be multithreaded.												

td_ta_new(3C_DB)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `libc_db(3LIB)`, `proc_service(3PROC)`, `attributes(5)`

td_ta_setconcurrency(3C_DB)

NAME td_ta_setconcurrency – set concurrency level for target process

SYNOPSIS

```
cc [ flag... ] file... -llibc_db [ library... ]
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_ta_setconcurrency(const td_thragent_t *ta_p, int level);
```

DESCRIPTION The `td_ta_setconcurrency()` function sets the desired concurrency level for the process identified by `ta_p` to level, just as if a thread within the process had called `thr_setconcurrency(3C)`.

RETURN VALUES

TD_OK	The call completed successfully.
TD_BADTA	An invalid internal process handle was passed in.
TD_BADPH	There is a NULL external process handle associated with this internal process handle. TD_NOCAPAB The client did not implement the <code>ps_kill(3PROC)</code> function in the imported interface.
TD_DBERR	A call to one of the imported interface routines failed.
TD_ERR	A <code>libc_db</code> internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `libc_db(3LIB)`, `ps_kill(3PROC)`, `thr_setconcurrency(3C)`, `attributes(5)`

td_ta_sync_iter(3C_DB)

NAME	td_ta_sync_iter, td_ta_thr_iter, td_ta_tsd_iter – iterator functions on process handles from libc_db
SYNOPSIS	<pre>cc [flag...] file... -lc_db [library...] #include <proc_service.h> #include <thread_db.h> typedef int td_sync_iter_f(const td_synchandle_t *sh_p, void *cbdata_p); typedef int td_thr_iter_f(const td_thrhandle_t *th_p, void *cbdata_p); typedef int td_key_iter_f(thread_key_t key, void (*destructor)(), void *cbdata_p); td_err_e td_ta_sync_iter(const td_thragent_t *ta_p, td_sync_iter_f *cb, void *cbdata_p); td_err_e td_ta_thr_iter(const td_thragent_t *ta_p, td_thr_iter_f *cb, void *cbdata_p, td_thr_state_e state, int ti_pri, sigset_t *ti_sigmask_p, unsigned ti_user_flags); td_err_e td_ta_tsd_iter(const td_thragent_t *ta_p, td_key_iter_f *cb, void *cbdata_p);</pre>
DESCRIPTION	<p>The <code>td_ta_sync_iter()</code>, <code>td_ta_thr_iter()</code>, and <code>td_ta_tsd_iter()</code> functions are iterator functions that when given a target process handle as an argument, return sets of handles for objects associated with the target process. The method is to call back a client-provided function once for each associated object, passing back a handle as well as the client-provided pointer <code>cb_data_p</code>. This enables a client to easily build a linked list of the associated objects. If the client-provided function returns non-zero, the iteration terminates, even if there are members remaining in the set of associated objects.</p> <p>The <code>td_ta_sync_iter()</code> function returns handles of synchronization objects (mutexes, readers-writer locks, semaphores, and condition variables) associated with a process. Some synchronization objects might not be known to <code>libc_db</code> and will not be returned. If the process has initialized the synchronization object (by calling <code>mutex_init(3C)</code>, for example) or a thread in the process has called a synchronization primitive (<code>mutex_lock()</code>, for example) using this object after <code>td_ta_new(3C_DB)</code> was called to attach to the process and <code>td_ta_sync_tracking_enable()</code> was called to enable synchronization object tracking, then a handle for the synchronization object will be passed to the callback function. See <code>td_sync_get_info(3C_DB)</code> for operations that can be performed on synchronization object handles.</p> <p>The <code>td_ta_thr_iter()</code> function returns handles for threads that are part of the target process. For <code>td_ta_thr_iter()</code>, the caller specifies several criteria to select a subset of threads for which the callback function should be called. Any of these selection criteria may be wild-carded. If all of them are wild-carded, then handles for all threads in the process will be returned.</p> <p>The selection parameters and corresponding wild-card values are:</p>

td_ta_sync_iter(3C_DB)

state (TD_THR_ANY_STATE):

Select only threads whose state matches state. See [td_thr_get_info\(3C_DB\)](#) for a list of thread states.

ti_pri (TD_THR_LOWEST_PRIORITY):

Select only threads for which the priority is at least ti_pri.

ti_sigmask_p (TD_SIGNO_MASK):

Select only threads whose signal mask exactly matches *ti_sigmask_p.

ti_user_flags (TD_THR_ANY_USER_FLAGS):

Select only threads whose user flags (specified at thread creation time) exactly match ti_user_flags.

The [td_ta_tsd_iter\(\)](#) function returns the thread-specific data keys in use by the current process. Thread-specific data for a particular thread and key can be obtained by calling [td_thr_tsd\(3C_DB\)](#).

RETURN VALUES

TD_OK The call completed successfully.
TD_BADTA An invalid process handle was passed.
TD_DBERR A call to one of the imported interface routines failed.
TD_ERR The call did not complete successfully.

ATTRIBUTES

See [attributes\(5\)](#) for description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO

[libc_db\(3LIB\)](#), [mutex_init\(3C\)](#), [td_sync_get_info\(3C_DB\)](#), [td_thr_get_info\(3C_DB\)](#), [td_thr_tsd\(3C_DB\)](#), [attributes\(5\)](#)

td_thr_dbsuspend(3C_DB)

NAME | td_thr_dbsuspend, td_thr_dbresume – suspend and resume threads in libc_db

SYNOPSIS |

```
cc [ flag... ] file... -lc_db [ library... ]
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_thr_dbsuspend(const td_thrhandle_t *th_p);
td_err_e td_thr_dbresume(const td_thrhandle_t *th_p);
```

DESCRIPTION | These operations do nothing other than call `ps_lstop(3PROC)` and `ps_lcontinue(3PROC)`, respectively, on the lightweight process (LWP) identified by the thread handle, *th_p*. Since `ps_lstop()` and `ps_lcontinue()` must be provided by the caller's application (see `proc_service(3PROC)`), and the application (a debugger-like entity) has full control over the stopped state of the process and all of its LWPs, `td_thr_dbsuspend()` and `td_thr_dbresume()` are unnecessary interfaces. They exist only to maintain interface compatibility with the past.

RETURN VALUES |

TD_OK	The call completed successfully.
TD_BADTH	An invalid thread handle was passed in.
TD_DBERR	A call to <code>ps_lstop()</code> or <code>ps_lcontinue()</code> failed.
TD_ERR	A <code>libc_db</code> internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO | `libc_db(3LIB)`, `attributes(5)`

NAME	td_thr_getgregs, td_thr_setgregs, td_thr_getfpregs, td_thr_setfpregs, td_thr_getxregsize, td_thr_getxregs, td_thr_setxregs – reading and writing thread registers in libc_db
SYNOPSIS	<pre>cc [flag...] file... -lc_db [library...] #include <proc_service.h> #include <thread_db.h> td_err_e td_thr_getgregs(const td_thrhandle_t *th_p, prgregset_t gregset); td_err_e td_thr_setgregs(const td_thrhandle_t *th_p, prgregset_t gregset); td_err_e td_thr_getfpregs(const td_thrhandle_t *th_p, prfpregset_t *fpregset); td_err_e td_thr_setfpregs(const td_thrhandle_t *th_p, prfpregset_t *fpregset); td_err_e td_thr_getxregsize(const td_thrhandle_t *th_p, int *xregsize); td_err_e td_thr_getxregs(const td_thrhandle_t *th_p, prxregset_t *xregset); td_err_e td_thr_setxregs(const td_thrhandle_t *th_p, prxregset_t *xregset);</pre>
DESCRIPTION	<p>These functions read and write the register sets associated with thread <i>th_p</i>. The <code>td_thr_getgregs()</code> and <code>td_thr_setgregs()</code> functions get and set, respectively, the general registers of thread <i>th_p</i>. The <code>td_thr_getfpregs()</code> and <code>td_thr_setfpregs()</code> functions get and set, respectively, the thread's floating point register set. The <code>td_thr_getxregsize()</code>, <code>td_thr_getxregs()</code>, and <code>td_thr_setxregs()</code> functions are SPARC-specific. The <code>td_thr_getxregsize()</code> function returns in <i>xregsize</i> the size of the architecture-dependent extra state registers. The <code>td_thr_getxregs()</code> and <code>td_thr_setxregs()</code> functions get and set, respectively, those extra state registers. On non-SPARC architectures, these functions return <code>TD_NOXREGS</code>.</p> <p>If the thread specified by <i>th_p</i> is currently executing on a lightweight process (LWP), these functions read or write, respectively, the appropriate register set to the LWP using the imported interface. If the thread is not currently executing on an LWP, the floating point and extra state registers may cannot be read or written. Some of the general registers might also not be readable or writable, depending on the architecture, in which case <code>td_thr_getfpregs()</code> and <code>td_thr_setfpregs()</code> return <code>TD_NOFPREGS</code> and <code>td_thr_getxregs()</code> and <code>td_thr_setxregs()</code> will return <code>TD_NOXREGS</code>. Calls to <code>td_thr_getgregs()</code> and <code>td_thr_setgregs()</code> succeed, but values returned for unreadable registers are undefined, values specified for unwritable registers are ignored. In this instance, and <code>TD_PARTIALREGS</code> is returned. See the architecture-specific notes that follow regarding the registers that may be read and written for a thread not currently executing on an LWP.</p>

td_thr_getgregs(3C_DB)

SPARC On a thread not currently assigned to an LWP, only %i0-%i7, %l0-%l7, %g7, %pc, and %sp (%o6) can be read or written. %pc and %sp refer to the program counter and stack pointer that the thread will have when it resumes execution.

x86 Architecture On a thread not currently assigned to an LWP, only %pc, %sp, %ebp, %edi, %edi, and %ebx can be read.

RETURN VALUES

TD_OK	The call completed successfully.
TD_BADTH	An invalid thread handle was passed in.
TD_DBERR	A call to one of the imported interface routines failed.
TD_PARTIALREGS	Because the thread is not currently assigned to a LWP, not all registers were read or written. See DESCRIPTION for a discussion about which registers are not saved when a thread is not assigned to an LWP.
TD_NOFPREGS	Floating point registers could not be read or written, either because the thread is not currently assigned to an LWP, or because the architecture does not have such registers.
TD_NOXREGS	Architecture-dependent extra state registers could not be read or written, either because the thread is not currently assigned to an LWP, or because the architecture does not have such registers, or because the architecture is not a SPARC architecture.
TD_ERR	A libc_db internal error occurred.

ATTRIBUTES See attributes(5) for description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO libc_db(3LIB), attributes(5)

NAME	td_thr_get_info – get thread information in libc_db library of interfaces
SYNOPSIS	<pre>cc [flag...] file... -lc_db [library...] #include <proc_service.h> #include <thread_db.h> td_err_e td_thr_get_info(const td_thrhandle_t *th_p, td_thrinfo_t *ti_p);</pre>
DESCRIPTION	<p>The <code>td_thr_get_info()</code> function fills in the <code>td_thrinfo_t</code> structure <code>*ti_p</code> with values for the thread identified by <code>th_p</code>.</p> <p>The <code>td_thrinfo_t</code> structure contains the following fields:</p> <pre>typedef struct td_thrinfo_t { td_thragen_tx *ti_ta_p /* internal process handle */ unsigned ti_user_flags; /* value of flags parameter */ thread_t ti_tid; /* thread identifier */ char *ti_tls; /* pointer to thread-local storage*/ paddr ti_startfunc; /* address of function at which thread execution began*/ paddr ti_stkbase; /* base of thread's stack area*/ int ti_stksize; /* size in bytes of thread's allocated stack region*/ paddr ti_ro_area; /* address of ulwp_t structure*/ int ti_ro_size /* size of the ulwp_t structure in bytes */ td_thr_state_e ti_state /* state of the thread */ uchar_t ti_db_suspended /* non-zero if thread suspended by td_thr_dbsuspend*/ td_thr_type_e ti_type /* type of the thread*/ int ti_pc /* value of thread's program counter*/ int ti_sp /* value of thread's stack counter*/ short ti_flags /* set of special flags used by libc*/ int ti_pri /* priority of thread returned by thr_getprio(3T)*/ lwpid_t ti_lid /* id of light weight process (LWP) executing this thread*/ sigset_t ti_sigmask /* thread's signal mask. See thr_sigsetmask(3T)*/ u_char ti_traceme /* non-zero if event tracing is on*/ uchar_t ti_preemptflag /* non-zero if thread preempted when last active*/ uchar_t ti_pirecflag /* non-zero if thread runs priority beside regular */ sigset_t ti_pending /* set of signals pending for this thread*/ td_thr_events_t ti_events /* bitmap of events enabled for this thread*/ };</pre> <p>The <code>ti_ta_p</code> member is the internal process handle identifying the process of which the thread is a member.</p> <p>The <code>ti_user_flags</code> member is the value of the flags parameter passed to <code>thr_create(3C)</code> when the thread was created.</p>

td_thr_get_info(3C_DB)

The `ti_tid` member is the thread identifier for the thread returned by `thr_create(3C)`.

The `ti_tls` member is the thread's pointer to thread-local storage.

The `ti_startfunc` member is the address of the function at which thread execution began, as specified when the thread was created with `thr_create(3C)`.

The `ti_stkbase` member is the base of the thread's stack area.

The `ti_stksize` member is the size in bytes of the thread's allocated stack region.

The `ti_ro_area` member is the address of the `ulwp_t` structure for this thread. Since accessing the `ulwp_t` structure directly violates the encapsulation provided by `libc_db`, this member should generally not be used. However, it might be useful as a prototype for extensions.

The `ti_state` member is the state of the thread. The `td_thr_state_e` enumeration type can contain the following values:

<code>TD_THR_ANY_STATE</code>	This value is never returned by <code>td_thr_get_info()</code> but is used as a wildcard to select threads in <code>td_ta_thr_iter()</code> .
<code>TD_THR_UNKNOWN</code>	The <code>libc_db</code> library cannot determine the state of the thread.
<code>TD_THR_STOPPED</code>	The thread has been stopped by a call to <code>thr_suspend(3C)</code> .
<code>TD_THR_RUN</code>	The thread is runnable, but it is not currently assigned to an LWP.
<code>TD_THR_ACTIVE</code>	The thread is currently executing on an LWP.
<code>TD_THR_ZOMBIE</code>	The thread has exited, but it has not yet been deallocated by a call to <code>thr_join(3C)</code> .
<code>TD_THR_SLEEP</code>	The thread is not currently runnable.
<code>TD_THR_STOPPED_ASLEEP</code>	The thread is both blocked by <code>TD_THR_SLEEP</code> and stopped by a call to <code>td_thr_dbsuspend(3C_DB)</code> .

The `ti_db_suspended` member is non-zero if and only if this thread is currently suspended because the controlling process has called `td_thr_dbsuspend` on it.

The `ti_type` member is a type of thread. It is either `TD_THR_USER` for a user thread (one created by the application), or `TD_THR_SYSTEM` for one created by `libc`.

td_thr_get_info(3C_DB)

The `ti_pc` member is the value of the thread's program counter, provided that the thread's `ti_state` value is `TD_THR_SLEEP`, `TD_THR_STOPPED`, or `TD_THR_STOPPED_ASLEEP`. Otherwise, the value of this member is undefined.

The `ti_sp` member is the value of the thread's stack pointer, provided that the thread's `ti_state` value is `TD_THR_SLEEP`, `TD_THR_STOPPED`, or `TD_THR_STOPPED_ASLEEP`. Otherwise, the value of this member is undefined.

The `ti_flags` member is a set of special flags used by `libc`, currently of use only to those debugging `libc`.

The `ti_pri` member is the thread's priority as it would be returned by `thr_getprio(3C)`.

The `ti_lid` member is the ID of the LWP executing this thread, or the ID of the LWP that last executed this thread, if this thread is not currently assigned to an LWP.

The `ti_sigmask` member is this thread's signal mask. See `thr_sigsetmask(3C)`.

The `ti_traceme` member is non-zero if and only if event tracing for this thread is on.

The `ti_preemptflag` member is non-zero if and only if the thread was preempted the last time it was active.

The `ti_pirecflag` member is non-zero if and only if due to priority inheritance the thread is currently running at a priority other than its regular priority.

The `ti_events` member is the bitmap of events enabled for this thread.

RETURN VALUES

- `TD_OK` The call completed successfully.
- `TD_BADTH` An invalid thread handle was passed in.
- `TD_DBERR` A call to one of the imported interface routines failed.
- `TD_ERR` The call did not complete successfully.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO

`libc_db(3LIB)`, `td_ta_thr_iter(3C_DB)`, `td_thr_dbsuspend(3C_DB)`, `thr_create(3C)`, `thr_getprio(3C)`, `thr_join(3C)`, `thr_sigsetmask(3C)`, `thr_suspend(3C)`, `attributes(5)`, `threads(5)`

td_thr_lockowner(3C_DB)

NAME | td_thr_lockowner – iterate over the set of locks owned by a thread

SYNOPSIS |

```
cc [ flag... ] file... -lc_db [ library... ]
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_thr_lockowner(const td_thrhandle_t *th_p,
    td_sync_iter_f *cb, void *cb_data_p);
```

DESCRIPTION | The `td_thr_lockowner()` function calls the iterator function *cb* once for every mutex that is held by the thread whose handle is *th_p*. The synchronization handle and the pointer *cb_data_p* are passed to the function. See `td_ta_thr_iter(3C_DB)` for a similarly structured function.

Iteration terminates early if the callback function *cb* returns a non-zero value.

RETURN VALUES |

TD_OK	The call completed successfully.
TD_BADTH	An invalid thread handle was passed in.
TD_BADPH	There is a NULL external process handle associated with this internal process handle.
TD_DBERR	A call to one of the imported interface routines failed.
TD_ERR	A <code>libc_db</code> internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO | `libc_db(3LIB)`, `td_ta_thr_iter(3C_DB)`, `attributes(5)`

NAME td_thr_setprio – set the priority of a thread

SYNOPSIS

```
cc [ flag... ] file... -lc_db [ library... ]
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_thr_setprio(const td_thrhandle_t *th_p, const int
    new_prio);
```

DESCRIPTION The `td_thr_setprio()` function sets priority of the thread specified by *th_p* to *new_prio*, as though a thread within the process had called `thr_setprio(3C)`.

RETURN VALUES

TD_OK	The call completed successfully.
TD_BADTH	An invalid thread handle was passed in.
TD_DBERR	A call to one of the imported interface routines failed.
TD_ERR	<i>new_prio</i> is an illegal value (out of range).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `libc_db(3LIB)`, `thr_setprio(3C)`, `attributes(5)`

td_thr_setsigpending(3C_DB)

NAME	td_thr_setsigpending, td_thr_sigsetmask – manage thread signals for libc_db								
SYNOPSIS	<pre>cc [flag...] file... -lc_db [library...] #include <proc_service.h> #include <thread_db.h> td_err_e td_thr_setsigpending(const td_thrhandle_t * th_p, const uchar_t ti_sigpending_flag, const sigset_t ti_sigmask); td_err_e td_thr_sigsetmask(const td_thrhandle_t * th_p, const sigset_t ti_sigmask);</pre>								
DESCRIPTION	<p>The <code>td_thr_setsigpending()</code> and <code>td_thr_sigsetmask()</code> functions affect the signal state of the thread identified by <code>th_p</code>.</p> <p>The <code>td_thr_setsigpending()</code> function sets the set of pending signals for thread <code>th_p</code> to <code>ti_sigpending</code>. The value of the libc-internal field that indicates whether a thread has any signal pending is set to <code>ti_sigpending_flag</code>. To be consistent, <code>ti_sigpending_flag</code> should be 0 if and only if all of the bits in <code>ti_sigpending</code> are 0.</p> <p>The <code>td_thr_sigsetmask()</code> function sets the signal mask of the thread <code>th_p</code> as if the thread had set its own signal mask with <code>thr_sigsetmask(3C)</code>. The new signal mask is the value of <code>ti_sigmask</code>.</p> <p>There is no equivalent to the <code>SIG_BLOCK</code> or <code>SIG_UNBLOCK</code> operations of <code>thr_sigsetmask(3C)</code>, which mask or unmask specific signals without affecting the mask state of other signals. To block or unblock specific signals,</p> <ol style="list-style-type: none"> 1. stop either the entire process or the thread with <code>td_thr_dbsuspend()</code>, 2. determine the thread's existing signal mask by calling <code>td_thr_get_info(3C_DB)</code>, 3. modify the <code>ti_sigmask</code> member of the <code>td_thrinfo_t</code> structure as desired, and 4. set the new signal mask with <code>td_thr_sigsetmask()</code>. 								
RETURN VALUES	<table border="0"> <tr> <td style="padding-right: 20px;">TD_OK</td> <td>The call completed successfully.</td> </tr> <tr> <td>TD_BADTH</td> <td>An invalid thread handle was passed in.</td> </tr> <tr> <td>TD_DBERR</td> <td>A call to one of the imported interface routines failed.</td> </tr> <tr> <td>TD_ERR</td> <td>A libc_db internal error occurred.</td> </tr> </table>	TD_OK	The call completed successfully.	TD_BADTH	An invalid thread handle was passed in.	TD_DBERR	A call to one of the imported interface routines failed.	TD_ERR	A libc_db internal error occurred.
TD_OK	The call completed successfully.								
TD_BADTH	An invalid thread handle was passed in.								
TD_DBERR	A call to one of the imported interface routines failed.								
TD_ERR	A libc_db internal error occurred.								
ATTRIBUTES	See <code>attributes(5)</code> for description of the following attributes:								
	<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>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe				
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
MT-Level	Safe								
SEE ALSO	<code>libc_db(3LIB)</code> , <code>td_thr_dbsuspend(3C_DB)</code> , <code>td_thr_get_info(3C_DB)</code> , <code>attributes(5)</code>								

td_thr_sleepinfo(3C_DB)

NAME td_thr_sleepinfo – return the synchronization handle for the object on which a thread is blocked

SYNOPSIS

```
cc [ flag... ] file... -lc_db [ library... ]
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_thr_sleepinfo(const td_thrhandle_t *th_p,
    td_synchandle_t *sh_p);
```

DESCRIPTION The td_thr_sleepinfo() function returns in *sh_p the handle of the synchronization object on which a sleeping thread is blocked.

RETURN VALUES

TD_OK	The call completed successfully.
TD_BADTH	An invalid thread handle was passed in.
TD_DBERR	A call to one of the imported interface routines failed.
TD_ERR	The thread th_p is not blocked on a synchronization object, or a libc_db internal error occurred.

ATTRIBUTES See attributes(5) for description of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO libc_db(3LIB), attributes(5)

td_thr_tsd(3C_DB)

NAME | td_thr_tsd – get a thread’s thread-specific data for libc_db library of interfaces

SYNOPSIS |

```
cc [ flag... ] file... -lc_db [ library... ]
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_thr_tsd(const td_thrhandle_t, const thread_key_t key,
void *data_pp);
```

DESCRIPTION | The `td_thr_tsd()` function returns in `*data_pp` the thread-specific data pointer for the thread identified by `th_p` and the thread-specific data key `key`. This is the same value that the thread `th_p` would obtain if it called `thr_getspecific(3C)`.

To find all the thread-specific data keys in use in a given target process, call `td_ta_tsd_iter(3C_DB)`.

RETURN VALUES |

TD_OK	The call completed successfully.
TD_BADTH	An invalid thread handle was passed in.
TD_DBERR	A call to one of the imported interface routines failed.
TD_ERR	A libc_db internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO | `libc_db(3LIB)`, `td_ta_tsd_iter(3C_DB)`, `thr_getspecific(3C)`, `attributes(5)`

NAME td_thr_validate – test a thread handle for validity

SYNOPSIS

```
cc [ flag... ] file... -lc_db [ library... ]
#include <proc_service.h>
#include <thread_db.h>

td_err_e td_thr_validate(const td_thrhandle_t *th_p);
```

DESCRIPTION The `td_thr_validate()` function tests whether `th_p` is a valid thread handle. A valid thread handle can become invalid if its thread exits.

RETURN VALUES

TD_OK	The call completed successfully. <code>th_p</code> is a valid thread handle.
TD_BADTH	<code>th_p</code> was NULL.
TD_DBERR	A call to one of the imported interface routines failed.
TD_NOTHR	<code>th_p</code> is not a valid thread handle.
TD_ERR	A <code>libc_db</code> internal error occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO `libc_db(3LIB)`, `attributes(5)`

tell(3C)

NAME tell – return a file offset for a file descriptor

SYNOPSIS #include <unistd.h>
off_t tell (int *fd*);

DESCRIPTION The tell() function obtains the current value of the file-position indicator for the file descriptor *fd*.

RETURN VALUES Upon successful completion, tell() returns the current value of the file-position indicator for *fd* measured in bytes from the beginning of the file.
Otherwise, it returns -1 and sets errno to indicate the error.

ERRORS The tell() function will fail if:

EBADF The file descriptor *fd* is not an open file descriptor.

E_OVERFLOW The current file offset cannot be represented correctly in an object of type off_t.

ESPIPE The file descriptor *fd* is associated with a pipe or FIFO.

USAGE The tell() function is equivalent to lseek(*fd*, 0, SEEK_CUR).

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO lseek(2), attributes(5)

NAME	telldir – current location of a named directory stream						
SYNOPSIS	<pre>#include <dirent.h> long int telldir(DIR *dirp);</pre>						
DESCRIPTION	<p>The <code>telldir()</code> function obtains the current location associated with the directory stream specified by <i>dirp</i>.</p> <p>If the most recent operation on the directory stream was a seekdir(3C), the directory position returned from the <code>telldir()</code> is the same as that supplied as a <i>loc</i> argument for <code>seekdir()</code>.</p>						
RETURN VALUES	Upon successful completion, <code>telldir()</code> returns the current location of the specified directory stream.						
ERRORS	<p>The <code>telldir()</code> function will fail if:</p> <p>EOVERFLOW The current location of the directory cannot be stored in an object of type <code>long</code>.</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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	opendir(3C) , readdir(3C) , seekdir(3C) , attributes(5) , standards(5)						

termios(3C)

NAME	termios – general terminal interface
SYNOPSIS	<pre>#include <termios.h> int tcgetattr(int <i>fildev</i>, struct termios <i>*termios_p</i>); int tcsetattr(int <i>fildev</i>, int <i>optional_actions</i>, const struct termios <i>*termios_p</i>); int tcsendbreak(int <i>fildev</i>, int <i>duration</i>); int tcdrain(int <i>fildev</i>); int tcflush(int <i>fildev</i>, int <i>queue_selector</i>); int tcflow(int <i>fildev</i>, int <i>action</i>); speed_t cfgetospeed(const struct termios <i>*termios_p</i>); int cfsetospeed(struct termios <i>*termios_p</i>, speed_t <i>speed</i>); speed_t cfgetispeed(const struct termios <i>*termios_p</i>); int cfsetispeed(struct termios <i>*termios_p</i>, speed_t <i>speed</i>); #include <sys/types.h> #include <termios.h> pid_t tcgetpgrp(int <i>fildev</i>); int tcsetpgrp(int <i>fildev</i>, pid_t <i>pgid</i>); pid_t tcgetsid(int <i>fildev</i>);</pre>
DESCRIPTION	<p>These functions describe a general terminal interface for controlling asynchronous communications ports. A more detailed overview of the terminal interface can be found in termio(7I), which also describes an ioctl(2) interface that provides the same functionality. However, the function interface described by these functions is the preferred user interface.</p> <p>Each of these functions is now described on a separate manual page.</p>
SEE ALSO	ioctl(2) , cfgetispeed(3C) , cfgetospeed(3C) , cfsetispeed(3C) , cfsetospeed(3C) , tcdrain(3C) , tcflow(3C) , tcflush(3C) , tcgetattr(3C) , tcgetpgrp(3C) , tcgetsid(3C) , tcsendbreak(3C) , tcsetattr(3C) , tcsetpgrp(3C) , tcsendbreak(3C) , termio(7I)

NAME	thr_create – create a thread				
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i>[<i>library...</i>] #include <thread.h> int thr_create(void *<i>stack_base</i>, size_t <i>stack_size</i>, void *(*<i>start_func</i>) (void*), void *<i>arg</i>, long <i>flags</i>, thread_t *<i>new_thread_ID</i>);</pre>				
DESCRIPTION	<p>Thread creation adds a new thread of control to the current process. The procedure <code>main()</code> is a single thread of control. Each thread executes concurrently with all other threads within the calling process and with other threads from other active processes.</p> <p>Although a newly created thread shares all of the calling process's global data with the other threads in the process, it has its own set of attributes and private execution stack. The new thread inherits the calling thread's signal mask and scheduling priority. Pending signals for a new thread are not inherited and will be empty.</p> <p>The call to create a thread takes the address of a user-defined function, specified by <i>start_func</i>, as one of its arguments. This function is the complete execution routine for the new thread.</p> <p>The lifetime of a thread begins with the successful return from <code>thr_create()</code>, which calls <i>start_func()</i> and ends with one of the following:</p> <ul style="list-style-type: none"> ■ the normal completion of <i>start_func()</i>, ■ the return from an explicit call to <code>thr_exit(3C)</code>, or ■ the conclusion of the calling process (see <code>exit(2)</code>). <p>The new thread performs by calling the function defined by <i>start_func</i> with only one argument, <i>arg</i>. If more than one argument needs to be passed to <i>start_func</i>, the arguments can be packed into a structure, the address of which can be passed to <i>arg</i>.</p> <p>If <i>start_func</i> returns, the thread terminates with the exit status set to the <i>start_func</i> return value (see <code>thr_exit(3C)</code>).</p> <p>When the thread from which <code>main()</code> originated returns, the effect is the same as if an implicit call to <code>exit()</code> were made using the return value of <code>main()</code> as the exit status. This behavior differs from a <i>start_func</i> return. If <code>main()</code> calls <code>thr_exit(3C)</code>, only the main thread exits, not the entire process.</p> <p>If the thread creation fails, a new thread is not created and the contents of the location referenced by the pointer to the new thread are undefined.</p> <p>The <i>flags</i> argument specifies which attributes are modifiable for the created thread. The value in <i>flags</i> is determined by the bitwise inclusive-OR of the following:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;">THR_BOUND</td> <td>This flag affects the contentionscope attribute of the thread. The new thread is created permanently bound to an LWP (that is, it is a <i>bound thread</i>.)</td> </tr> <tr> <td>THR_DETACHED</td> <td>This flag affects the detachstate attribute of the thread. The new thread is created detached. The exit status of a detached thread is</td> </tr> </table>	THR_BOUND	This flag affects the contentionscope attribute of the thread. The new thread is created permanently bound to an LWP (that is, it is a <i>bound thread</i> .)	THR_DETACHED	This flag affects the detachstate attribute of the thread. The new thread is created detached. The exit status of a detached thread is
THR_BOUND	This flag affects the contentionscope attribute of the thread. The new thread is created permanently bound to an LWP (that is, it is a <i>bound thread</i> .)				
THR_DETACHED	This flag affects the detachstate attribute of the thread. The new thread is created detached. The exit status of a detached thread is				

thr_create(3C)

not accessible to other threads. Its thread ID and other resources may be re-used as soon as the thread terminates. [thr_join\(3C\)](#) will not wait for a detached thread.

THR_NEW_LWP	This flag is obsolete and is maintained for compatibility.
THR_SUSPENDED	This flag affects the suspended attribute of the thread. The new thread is created suspended and will not execute <i>start_func</i> until it is started by <code>thr_continue()</code> .
THR_DAEMON	This flag affects the daemon attribute of the thread. In addition to being created detached (THR_DAEMON implies THR_DETACHED), the thread is marked as a daemon. Daemon threads do not interfere with the exit conditions for a process. A process will terminate when the last non-daemon thread exits or the process calls <code>exit(2)</code> . Also, a thread that is waiting in thr_join(3C) for any thread to terminate will return EDEADLK when all remaining threads in the process are either daemon threads or other threads waiting in <code>thr_join()</code> . Daemon threads are most useful in libraries that want to use threads.

Default thread creation:

```
thread_t tid;
void *start_func(void *), *arg;
thr_create(NULL, NULL, start_func, arg, NULL, &tid);
```

User-defined thread creation (create a bound permanently to an LWP, that is, a bound thread):

```
thr_create(NULL, NULL, start_func, arg, THR_BOUND, &tid);
```

With `thr_create()`, the new thread uses the stack beginning at the address specified by *stack_base* and continuing for *stack_size* bytes. The *stack_size* argument must be greater than the value returned by [thr_min_stack\(3C\)](#). If *stack_base* is NULL, `thr_create()` allocates a stack for the new thread with at least *stack_size* bytes. If *stack_size* is 0, a default size is used. If *stack_size* is not 0, it must be greater than the value returned by [thr_min_stack\(3C\)](#) See NOTES.

When *new_thread_ID* is not NULL, it points to a location where the ID of the new thread is stored if `thr_create()` is successful. The ID is only valid within the calling process.

RETURN VALUES

If successful, the `thr_create()` function returns 0. Otherwise, an error value is returned to indicate the error. If the application is not linked with the threads library, -1 is returned.

ERRORS

EAGAIN	A resource control limit on the total number of threads in a process, task, project, or zone has been exceeded or some system resource has been exceeded (for example, too many LWPs were created).
--------	---

EINVAL	The <i>stack_base</i> argument is not NULL and <i>stack_size</i> is less than the value returned by <code>thr_min_stack(3C)</code> , or the <i>stack_base</i> argument is NULL and <i>stack_size</i> is not 0 and is less than the value returned by <code>thr_min_stack(3C)</code> .
ENOMEM	The system cannot allocate stack for the thread.

The `thr_create()` function may use `mmap()` to allocate thread stacks from `MAP_PRIVATE`, `MAP_NORESERVE`, and `MAP_ANON` memory mappings if *stack_base* is NULL, and consequently may return upon failure the relevant error values returned by `mmap()`. See the `mmap(2)` manual page for these error values.

EXAMPLES

The following is an example of concurrency with multithreading. Since POSIX threads and Solaris threads are fully compatible even within the same process, this example uses `pthread_create()` if you execute `a.out 0`, or `thr_create()` if you execute `a.out 1`.

Five threads are created that simultaneously perform a time-consuming function, `sleep(10)`. If the execution of this process is timed, the results will show that all five individual calls to `sleep` for ten-seconds completed in about ten seconds, even on a uniprocessor. If a single-threaded process calls `sleep(10)` five times, the execution time will be about 50-seconds.

The command-line to time this process is:

```
/usr/bin/time a.out 0 (for POSIX threading)
```

or

```
/usr/bin/time a.out 1 (for Solaris threading)
```

EXAMPLE 1 An example of concurrency with multithreading.

```
/* cc thisfile.c -lthread -lpthread */
#define _REENTRANT /* basic 3-lines for threads */
#include <pthread.h>
#include <thread.h>
#define NUM_THREADS 5
#define SLEEP_TIME 10

void *sleeping(void *); /* thread routine */
int i;
thread_t tid[NUM_THREADS]; /* array of thread IDs */

int
main(int argc, char *argv[])
{
    if (argc == 1) {
        printf("use 0 as arg1 to use pthread_create( )\n");
        printf("or use 1 as arg1 to use thr_create( )\n");
        return (1);
    }
}
```

thr_create(3C)

EXAMPLE 1 An example of concurrency with multithreading. (Continued)

```
switch (*argv[1]) {
case '0': /* POSIX */
    for ( i = 0; i < NUM_THREADS; i++)
        pthread_create(&tid[i], NULL, sleeping,
            (void *)SLEEP_TIME);
    for ( i = 0; i < NUM_THREADS; i++)
        pthread_join(tid[i], NULL);
    break;

case '1': /* Solaris */
    for ( i = 0; i < NUM_THREADS; i++)
        thr_create(NULL, 0, sleeping, (void *)SLEEP_TIME, 0,
            &tid[i]);
    while (thr_join(0, NULL, NULL) == 0)
        ;
    break;
} /* switch */
printf("main( ) reporting that all %d threads have terminated\n", i);
return (0);
} /* main */

void *
sleeping(void *arg)
{
    int sleep_time = (int)arg;
    printf("thread %d sleeping %d seconds ...\n", thr_self( ), sleep_time);
    sleep(sleep_time);
    printf("\nthread %d awakening\n", thr_self( ));
    return (NULL);
}
```

Had `main()` not waited for the completion of the other threads (using `pthread_join(3C)` or `thr_join(3C)`), it would have continued to process concurrently until it reached the end of its routine and the entire process would have exited prematurely (see `exit(2)`).

EXAMPLE 2 Creating a default thread with a new signal mask.

The following example demonstrates how to create a default thread with a new signal mask. The *new_mask* argument is assumed to have a value different from the creator's signal mask (*orig_mask*). The *new_mask* argument is set to block all signals except for `SIGINT`. The creator's signal mask is changed so that the new thread inherits a different mask, and is restored to its original value after `thr_create()` returns.

This example assumes that `SIGINT` is also unmasked in the creator. If it is masked by the creator, then unmasking the signal opens the creator to this signal. The other alternative is to have the new thread set its own signal mask in its start routine.

```
thread_t tid;
sigset_t new_mask, orig_mask;
int error;
```

EXAMPLE 2 Creating a default thread with a new signal mask. (Continued)

```
(void) sigfillset(&new_mask);
(void) sigdelset(&new_mask, SIGINT);
(void) thr_sigsetmask(SIG_SETMASK, &new_mask, &orig_mask);
error = thr_create(NULL, 0, do_func, NULL, 0, &tid);
(void) thr_sigsetmask(SIG_SETMASK, &orig_mask, NULL);
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `exit(2)`, `getrlimit(2)`, `mmap(2)`, `exit(3C)`, `sleep(3C)`, `thr_exit(3C)`, `thr_join(3C)`, `thr_min_stack(3C)`, `thr_setconcurrency(3C)`, `thr_suspend(3C)`, `attributes(5)`, `standards(5)`, `threads(5)`

NOTES Since multithreaded-application threads execute independently of each other, their relative behavior is unpredictable. It is therefore possible for the thread executing `main()` to finish before all other user-application threads.

Using `thr_join(3C)` in the following syntax,

```
while (thr_join(0, NULL, NULL) == 0);
```

will cause the invoking thread (which may be `main()`) to wait for the termination of all non-daemon threads, excluding threads that are themselves waiting in `thr_join()`; however, the second and third arguments to `thr_join()` need not necessarily be `NULL`.

A thread has not terminated until `thr_exit()` has finished. The only way to determine this is by `thr_join()`. When `thr_join()` returns a departed thread, it means that this thread has terminated and its resources are reclaimable. For instance, if a user specified a stack to `thr_create()`, this stack can only be reclaimed after `thr_join()` has reported this thread as a departed thread. It is not possible to determine when a *detached* thread has terminated. A detached thread disappears without leaving a trace.

Typically, thread stacks allocated by `thr_create()` begin on page boundaries and any specified (a red-zone) size is rounded up to the next page boundary. A page with no access permission is appended to the top of the stack so that most stack overflows will result in a `SIGSEGV` signal being sent to the offending thread. Thread stacks allocated by the caller are used as is.

Using a default stack size for the new thread, instead of passing a user-specified stack size, results in much better `thr_create()` performance. The default stack size for a user-thread is 1 megabyte in a 32-bit process and 2 megabyte in a 64-bit process.

thr_create(3C)

A user-specified stack size must be greater than the value `THR_MIN_STACK`. A minimum stack size may not accommodate the stack frame for the user thread function *start_func*. If a stack size is specified, it must accommodate *start_func* requirements and the functions that it may call in turn, in addition to the minimum requirement.

It is usually very difficult to determine the runtime stack requirements for a thread. `THR_MIN_STACK` specifies how much stack storage is required to execute a `NULL` *start_func*. The total runtime requirements for stack storage are dependent on the storage required to do runtime linking, the amount of storage required by library runtimes (like `printf()`) that your thread calls. Since these storage parameters are not known before the program runs, it is best to use default stacks. If you know your runtime requirements or decide to use stacks that are larger than the default, then it makes sense to specify your own stacks.

NAME	thr_exit – terminate the calling thread
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <thread.h> void thr_exit(void *status);</pre>
DESCRIPTION	<p>The <code>thr_exit()</code> function terminates the calling thread, in a similar way that <code>exit(3C)</code> terminates the calling process. If the calling thread is not detached, then the thread's ID and the exit status specified by <code>status</code> are retained. The value <code>status</code> is then made available to any successful join with the terminating thread (see <code>thr_join(3C)</code>); otherwise, <code>status</code> is disregarded allowing the thread's ID to be reclaimed immediately.</p> <p>Any cancellation cleanup handlers that have been pushed and not yet popped are popped in the reverse order that they were pushed and then executed. After all cancellation cleanup handlers have been executed, if the thread has any thread-specific data, appropriate destructor functions will be called in an unspecified order. Thread termination does not release any application visible process resources, including, but not limited to, mutexes and file descriptors, nor does it perform any process level cleanup actions, including, but not limited to, calling any <code>atexit()</code> routines that might exist.</p> <p>If any thread, including the <code>main()</code> thread, calls <code>thr_exit()</code>, only that thread will exit.</p> <p>If <code>main()</code> returns or exits (either implicitly or explicitly), or any thread explicitly calls <code>exit()</code>, the entire process will exit.</p> <p>The behavior of <code>thr_exit()</code> is undefined if called from a cancellation cleanup handler or destructor function that was invoked as a result of either an implicit or explicit call to <code>thr_exit()</code>.</p> <p>After a thread has terminated, the result of access to local (auto) variables of the thread is undefined. Thus, references to local variables of the exiting thread should not be used for the <code>thr_exit()</code> <code>status</code> parameter value.</p> <p>If any thread (except the <code>main()</code> thread) implicitly or explicitly returns, the result is the same as if the thread called <code>thr_exit()</code> and it will return the value of <code>status</code> as the exit code.</p> <p>The process will terminate with an exit status of 0 after the last non-daemon thread has terminated (including the <code>main()</code> thread). This behavior is the same as if the application had called <code>exit()</code> with a 0 argument at thread termination time.</p>
RETURN VALUES	The <code>thr_exit()</code> function cannot return to its caller.
ERRORS	No errors are defined.

thr_exit(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `exit(3C)`, `thr_create(3C)`, `thr_join(3C)`, `thr_keycreate(3C)`, `attributes(5)`, `standards(5)`

NOTES Although only POSIX implements cancellation, cancellation can be used with Solaris threads, due to their interoperability.

The *status* argument should not reference any variables local to the calling thread.

NAME	thr_getconcurrency, thr_setconcurrency – get or set thread concurrency level				
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <thread.h> int thr_setconcurrency(int <i>new_level</i>); int thr_getconcurrency(void);</pre>				
DESCRIPTION	<p>These functions are obsolete and maintained for compatibility only. The <code>thr_setconcurrency()</code> function updates the desired concurrency level that <code>libthread</code> maintains for the calling process. This value does not affect the behavior of the calling process.</p> <p>The <code>thr_getconcurrency()</code> function returns the current value for the desired concurrency level.</p>				
RETURN VALUES	<p>The <code>thr_getconcurrency()</code> function always returns the current value for the desired concurrency level.</p> <p>If successful, the <code>thr_setconcurrency()</code> function returns 0. Otherwise, a non-zero value is returned to indicate the error.</p>				
ERRORS	<p>The <code>thr_setconcurrency()</code> function will fail if:</p> <p><code>EAGAIN</code> The specified concurrency level would cause a system resource to be exceeded.</p> <p><code>EINVAL</code> The value for <i>new_level</i> is negative.</p>				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"> <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>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>thr_create(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>				

thr_getprio(3C)

NAME	thr_getprio, thr_setprio – access dynamic thread scheduling				
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <thread.h> int thr_setprio(thread_t <i>target_thread</i>, int <i>priority</i>); int thr_getprio(thread_t <i>target_thread</i>, int *<i>priority</i>);</pre>				
DESCRIPTION	<p>The <code>thr_setprio()</code> function dynamically changes the priority of the thread specified by <i>target_thread</i> within the current process to the priority specified by <i>priority</i>. By default, threads contend for synchronization objects based on fixed priorities that range from 0, the least significant, to 127. The <i>target_thread</i> will receive precedence by <code>libthread</code> over lower priority threads with respect to synchronization object contention.</p> <p>The <code>thr_getprio()</code> function stores the current priority for the thread specified by <i>target_thread</i> in the location pointed to by <i>priority</i>. Thread priorities regulate the order in which threads unblock from synchronization objects and are different from realtime priorities, which regulate and enforce access to CPU resources. Programs that need access to "real" priorities should use bound threads in the realtime class (see <code>prionctl(2)</code>).</p>				
RETURN VALUES	If successful, the <code>thr_getprio()</code> and <code>thr_setprio()</code> return 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>For each of the following conditions, these functions return an error number if the condition is detected.</p> <p>ESRCH The value specified by <i>target_thread</i> does not refer to an existing thread.</p> <p>The <code>thr_getprio()</code> and <code>thr_setprio()</code> functions may fail if:</p> <p>EINVAL The value of <i>priority</i> makes no sense for the scheduling class associated with the <i>target_thread</i>.</p>				
ATTRIBUTES	See <code>attributes(5)</code> 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>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>prionctl(2)</code> , <code>sched_setparam(3RT)</code> , <code>thr_create(3C)</code> , <code>thr_suspend(3C)</code> , <code>thr_yield(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>				

NAME	thr_join – wait for thread termination				
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <thread.h> int thr_join(thread_t <i>thread</i>, thread_t *<i>departed</i>, void **<i>status</i>);</pre>				
DESCRIPTION	<p>The <code>thr_join()</code> function suspends processing of the calling thread until the target <i>thread</i> completes. The <i>thread</i> argument must be a member of the current process and cannot be a detached thread. See thr_create(3C).</p> <p>If two or more threads wait for the same thread to complete, all will suspend processing until the thread has terminated, and then one thread will return successfully and the others will return with an error of ESRCH. The <code>thr_join()</code> function will not block processing of the calling thread if the target <i>thread</i> has already terminated.</p> <p>If a <code>thr_join()</code> call returns successfully with a non-null <i>status</i> argument, the value passed to thr_exit(3C) by the terminating thread will be placed in the location referenced by <i>status</i>.</p> <p>If the target <i>thread</i> ID is 0, <code>thr_join()</code> finds and returns the status of a terminated undetached thread in the process. If no such thread exists, it suspends processing of the calling thread until a thread for which no other thread is waiting enters that state, at which time it returns successfully, or until all other threads in the process are either daemon threads or threads waiting in <code>thr_join()</code>, in which case it returns EDEADLK. See NOTES.</p> <p>If <i>departed</i> is not NULL, it points to a location that is set to the ID of the terminated thread if <code>thr_join()</code> returns successfully.</p>				
RETURN VALUES	If successful, <code>thr_join()</code> returns 0. Otherwise, an error number is returned to indicate the error.				
ERRORS	<p>EDEADLK A joining deadlock would occur, such as when a thread attempts to wait for itself, or the calling thread is waiting for any thread to exit and only daemon threads or waiting threads exist in the process.</p> <p>ESRCH No undetached thread could be found corresponding to the given thread ID.</p>				
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:				
	<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>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	thr_create(3C) , thr_exit(3C) , wait(3C) , attributes(5) , standards(5)				
NOTES	Using thr_join(3C) in the following syntax,				

thr_join(3C)

```
while (thr_join(0, NULL, NULL) == 0);
```

will wait for the termination of all non-daemon threads, excluding threads that are themselves waiting in `thr_join()`.

NAME	thr_keycreate, thr_setspecific, thr_getspecific – thread-specific-data functions
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <thread.h> int thr_keycreate(thread_key_t *keyp, void (*<i>destructor</i>)(void *)); int thr_setspecific(thread_key_t key, void *<i>value</i>); int thr_getspecific(thread_key_t key, void **<i>valuep</i>);</pre>
DESCRIPTION	
Create Key	<p>In general, thread key creation allocates a key that locates data specific to each thread in the process. The key is global to all threads in the process, which allows each thread to bind a value to the key once the key has been created. The key independently maintains specific values for each binding thread. The <code>thr_keycreate()</code> function allocates a global <i>key</i> namespace, pointed to by <i>keyp</i>, that is visible to all threads in the process. Each thread is initially bound to a private element of this <i>key</i>, which allows access to its thread-specific data.</p> <p>Upon key creation, a new key is assigned the value <code>NULL</code> for all active threads. Additionally, upon thread creation, all previously created keys in the new thread are assigned the value <code>NULL</code>.</p> <p>Optionally, a destructor function <i>destructor</i> can be associated with each <i>key</i>. Upon thread exit, if a <i>key</i> has a non-null <i>destructor</i> function and the thread has a non-null <i>value</i> associated with that <i>key</i>, the <i>destructor</i> function is called with the current associated <i>value</i>. If more than one <i>destructor</i> exists for a thread when it exits, the order of destructor calls is unspecified.</p>
Set Value	<p>Once a key has been created, each thread can bind a new <i>value</i> to the key using <code>thr_setspecific()</code>. The values are unique to the binding thread and are individually maintained. These values continue for the life of the calling thread.</p> <p>Proper synchronization of <i>key</i> storage and access must be ensured by the caller. The <i>value</i> argument to <code>thr_setspecific()</code> is generally a pointer to a block of dynamically allocated memory reserved by the calling thread for its own use. See EXAMPLES below.</p> <p>At thread exit, the <i>destructor</i> function, which is associated at time of creation, is called and it uses the specific key value as its sole argument.</p>
Get Value	<code>thr_getspecific()</code> stores the current value bound to <i>key</i> for the calling thread into the location pointed to by <i>valuep</i> .
RETURN VALUES	If successful, <code>thr_keycreate()</code> , <code>thr_setspecific()</code> and <code>thr_getspecific()</code> return 0. Otherwise, an error number is returned to indicate the error.
ERRORS	If the following conditions occur, <code>thr_keycreate()</code> returns the corresponding error number:

thr_keycreate(3C)

EAGAIN The system lacked the necessary resources to create another thread-specific data key.

ENOMEM Insufficient memory exists to create the key.

If the following conditions occur, `thr_keycreate()` and `thr_setspecific()` return the corresponding error number:

ENOMEM Insufficient memory exists to associate the value with the key.

The `thr_setspecific()` function returns the corresponding error number:

EINVAL The *key* value is invalid.

EXAMPLES

EXAMPLE 1 Call the thread-specific data from more than one thread without special initialization.

In this example, the thread-specific data in this function can be called from more than one thread without special initialization. For each argument passed to the executable, a thread is created and privately bound to the string-value of that argument.

```
/* cc -mt thisfile.c */

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

void *thread_specific_data(void *);
void cleanup(void*);
#define MAX_ARGC 20
thread_t tid[MAX_ARGC];
int num_threads;

int
main(int argc, char *argv[]) {
    int i;
    num_threads = argc - 1;
    for (i = 0; i < num_threads; i++)
        thr_create(NULL, 0, thread_specific_data, argv[i+1], 0, &tid[i]);
    for (i = 0; i < num_threads; i++)
        thr_join(tid[i], NULL, NULL);
    return (0);
} /* end main */

void *
thread_specific_data(void *arg) {
    static mutex_t keylock; /* static ensures only one copy of keylock */
    static thread_key_t key;
    static int once_per_keyname = 0;
    char *private_data = arg;
    void *tsd = NULL;
    void *data;

    if (!once_per_keyname) {
        mutex_lock(&keylock);
```


EXAMPLE 1 Call the thread-specific data from more than one thread without special initialization. (Continued)

```

        if (!once_per_keyname) {
            thr_keycreate(&key, cleanup);
            once_per_keyname++;
        }
        mutex_unlock(&keylock);
    }
    thr_getspecific(key, &tsd);
    if (tsd == NULL) {
        data = malloc(strlen(private_data) + 1);
        strcpy(data, private_data);
        thr_setspecific(key, data);
        thr_getspecific(key, &tsd);
    }
    printf("tsd for %d = %s\n",
        thr_self(), (char *)tsd);
    thr_getspecific(key, &tsd);
    printf("tsd for %d remains %s\n",
        thr_self(), (char *)tsd);
    return (NULL);
} /* end thread_specific_data */

void
cleanup(void *v) {
    /* application-specific clean-up function */
    free(v);
}

```

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

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

SEE ALSO [thr_exit\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

WARNINGS The `thr_getspecific()` and `thr_setspecific()` functions can be called either explicitly or implicitly from a thread-specific data destructor function. Calling `thr_setspecific()` from a destructor can result in lost storage or infinite loops.

thr_kill(3C)

NAME thr_kill – send a signal to a thread

SYNOPSIS

```
cc -mt [ flag... ] file... [ library... ]

#include <signal.h>
#include <thread.h>

int thr_kill(thread_t thread, int sig);
```

DESCRIPTION The `thr_kill()` function sends the *sig* signal to the thread designated by *thread*. *thread* must be a member of the same process as the calling thread. The *sig* argument must be one of the signals listed in `signal.h(3HEAD)`; with the exception of `SIGLWP`, `SIGCANCEL`, and `SIGWAITING` being reserved and off limits to `thr_kill()`. If *sig* is 0, a validity check is done for the existence of the target thread; no signal is sent.

RETURN VALUES Upon successful completion, `thr_kill()` returns 0. Otherwise, an error number is returned. In the event of failure, no signal is sent.

ERRORS

ESRCH	No thread was found that corresponded to the thread designated by <i>thread</i> ID.
EINVAL	The <i>sig</i> argument value is not zero and is an invalid or an unsupported signal number.

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

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

SEE ALSO `kill(2)`, `sigaction(2)`, `raise(3C)`, `signal.h(3HEAD)`, `thr_self(3C)`, `attributes(5)`, `standards(5)`

NAME thr_main – identify the main thread

SYNOPSIS `cc -mt [flag...] file... [library...]`

```
#include <thread.h>
int thr_main(void);
```

DESCRIPTION The `thr_main()` function returns one of the following:

- 1 if the calling thread is the main thread
- 0 if the calling thread is not the main thread
- 1 if `libthread` is not linked in or thread initialization has not completed

FILES `/lib/libthread`

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [thr_self\(3C\)](#), [attributes\(5\)](#)

thr_min_stack(3C)

NAME	thr_min_stack – return the minimum-allowable size for a thread’s stack				
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <thread.h> size_t thr_min_stack(void);</pre>				
DESCRIPTION	<p>When a thread is created with a user-supplied stack, the user must reserve enough space to run this thread. In a dynamically linked execution environment, it is very hard to know what the minimum stack requirements are for a thread. The function <code>thr_min_stack()</code> returns the amount of space needed to execute a null thread. This is a thread that was created to execute a null procedure. A thread that does something useful should have a stack size that is <code>thr_min_stack() + <some increment></code>.</p> <p>Most users should not be creating threads with user-supplied stacks. This functionality was provided to support applications that wanted complete control over their execution environment.</p> <p>Typically, users should let the threads library manage stack allocation. The threads library provides default stacks which should meet the requirements of any created thread.</p> <p><code>thr_min_stack()</code> will return the unsigned int <code>THR_MIN_STACK</code>, which is the minimum-allowable size for a thread’s stack.</p> <p>In this implementation the default size for a user-thread’s stack is one mega-byte. If the second argument to <code>thr_create(3C)</code> is <code>NULL</code>, then the default stack size for the newly-created thread will be used. Otherwise, you may specify a stack-size that is at least <code>THR_MIN_STACK</code>, yet less than the size of your machine’s virtual memory.</p> <p>It is recommended that the default stack size be used.</p> <p>To determine the smallest-allowable size for a thread’s stack, execute the following:</p> <pre>/* cc thisfile.c -lthread */ #define _REENTRANT #include <thread.h> #include <stdio.h> main() { printf("thr_min_stack() returns %u\n", thr_min_stack()); }</pre>				
ATTRIBUTES	See <code>attributes(5)</code> 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>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>attributes(5)</code> , <code>standards(5)</code>				

NAME thr_self – get calling thread’s ID

SYNOPSIS `cc -mt [flag...] file...[library...]
#include <thread.h>`

```
thread_t thr_self(void);
typedef(unsigned int thread_t);
```

DESCRIPTION thr_self() returns the thread ID of the calling thread.

ERRORS No errors are defined.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [thr_create\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

thr_sigsetmask(3C)

NAME	thr_sigsetmask – change or examine calling thread’s signal mask						
SYNOPSIS	<pre>cc -mt [<i>flag...</i>] <i>file...</i> [<i>library...</i>] #include <thread.h> #include <signal.h> int thr_sigsetmask(int <i>how</i>, const sigset_t *<i>set</i>, sigset_t *<i>oset</i>);</pre>						
DESCRIPTION	<p>The <code>thr_sigsetmask()</code> function changes or examines a calling thread’s signal mask. Each thread has its own signal mask. A new thread inherits the calling thread’s signal mask and priority; however, pending signals are not inherited. Signals pending for a new thread will be empty.</p> <p>If the value of the argument <i>set</i> is not <code>NULL</code>, <i>set</i> points to a set of signals that can modify the currently blocked set. If the value of <i>set</i> is <code>NULL</code>, the value of <i>how</i> is insignificant and the thread’s signal mask is unmodified; thus, <code>thr_sigsetmask()</code> can be used to inquire about the currently blocked signals.</p> <p>The value of the argument <i>how</i> specifies the method in which the set is changed and takes one of the following values:</p> <table><tr><td><code>SIG_BLOCK</code></td><td><i>set</i> corresponds to a set of signals to block. They are added to the current signal mask.</td></tr><tr><td><code>SIG_UNBLOCK</code></td><td><i>set</i> corresponds to a set of signals to unblock. These signals are deleted from the current signal mask.</td></tr><tr><td><code>SIG_SETMASK</code></td><td><i>set</i> corresponds to the new signal mask. The current signal mask is replaced by <i>set</i>.</td></tr></table> <p>If the value of <i>oset</i> is not <code>NULL</code>, it points to the location where the previous signal mask is stored.</p>	<code>SIG_BLOCK</code>	<i>set</i> corresponds to a set of signals to block. They are added to the current signal mask.	<code>SIG_UNBLOCK</code>	<i>set</i> corresponds to a set of signals to unblock. These signals are deleted from the current signal mask.	<code>SIG_SETMASK</code>	<i>set</i> corresponds to the new signal mask. The current signal mask is replaced by <i>set</i> .
<code>SIG_BLOCK</code>	<i>set</i> corresponds to a set of signals to block. They are added to the current signal mask.						
<code>SIG_UNBLOCK</code>	<i>set</i> corresponds to a set of signals to unblock. These signals are deleted from the current signal mask.						
<code>SIG_SETMASK</code>	<i>set</i> corresponds to the new signal mask. The current signal mask is replaced by <i>set</i> .						
RETURN VALUES	Upon successful completion, the <code>thr_sigsetmask()</code> function returns 0. Otherwise, it returns a non-zero value.						
ERRORS	The <code>thr_sigsetmask()</code> function will fail if: <table><tr><td><code>EINVAL</code></td><td>The value of <i>how</i> is not defined and <i>oset</i> is <code>NULL</code>.</td></tr></table>	<code>EINVAL</code>	The value of <i>how</i> is not defined and <i>oset</i> is <code>NULL</code> .				
<code>EINVAL</code>	The value of <i>how</i> is not defined and <i>oset</i> is <code>NULL</code> .						
EXAMPLES	<p>EXAMPLE 1 Create a default thread that can serve as a signal catcher/handler with its own signal mask.</p> <p>The following example shows how to create a default thread that can serve as a signal catcher/handler with its own signal mask. <i>new</i> will have a different value from the creator’s signal mask.</p> <p>As POSIX threads and Solaris threads are fully compatible even within the same process, this example uses <code>pthread_create(3C)</code> if you execute <code>a.out 0</code>, or <code>thr_create(3C)</code> if you execute <code>a.out 1</code>.</p> <p>In this example:</p>						

EXAMPLE 1 Create a default thread that can serve as a signal catcher/handler with its own signal mask. (Continued)

- The `sigemptyset(3C)` function initializes a null signal set, `new`. The `sigaddset(3C)` function packs the signal, `SIGINT`, into that new set.
- Either `pthread_sigmask()` or `thr_sigsetmask()` is used to mask the signal, `SIGINT` (CTRL-C), from the calling thread, which is `main()`. The signal is masked to guarantee that only the new thread will receive this signal.
- `pthread_create()` or `thr_create()` creates the signal-handling thread.
- Using `pthread_join(3C)` or `thr_join(3C)`, `main()` then waits for the termination of that signal-handling thread, whose ID number is `user_threadID`. Then `main()` will `sleep(3C)` for 2 seconds, after which the program terminates.
- The signal-handling thread, handler:
 - Assigns the handler `interrupt()` to handle the signal `SIGINT` by the call to `sigaction(2)`.
 - Resets its own signal set to *not block* the signal, `SIGINT`.
 - Sleeps for 8 seconds to allow time for the user to deliver the signal `SIGINT` by pressing the CTRL-C.

```

/* cc thisfile.c -lthread -lpthread */
#define _REENTRANT /* basic first 3-lines for threads */
#include <pthread.h>
#include <thread.h>

thread_t user_threadID;
sigset_t new;
void *handler( ), interrupt( );

main( int argc, char *argv[ ] ){
    test_argv(argv[1]);

    sigemptyset(&new);
    sigaddset(&new, SIGINT);
    switch(*argv[1]) {

        case '0': /* POSIX */
            pthread_sigmask(SIG_BLOCK, &new, NULL);
            pthread_create(&user_threadID, NULL, handler, argv[1]);
            pthread_join(user_threadID, NULL);
            break;

        case '1': /* Solaris */
            thr_sigsetmask(SIG_BLOCK, &new, NULL);
            thr_create(NULL, 0, handler, argv[1], 0, &user_threadID);
            thr_join(user_threadID, NULL, NULL);
            break;
    } /* switch */

    printf("thread handler, # %d, has exited\n",user_threadID);
    sleep(2);

```

thr_sigsetmask(3C)

EXAMPLE 1 Create a default thread that can serve as a signal catcher/handler with its own signal mask. (Continued)

```
        printf("main thread, # %d is done\n", thr_self( ));
    } /* end main */

    struct sigaction act;

    void *
    handler(char argvl[ ])
    {
        act.sa_handler = interrupt;
        sigaction(SIGINT, &act, NULL);
        switch(*argvl){
            case '0': /* POSIX */
                pthread_sigmask(SIG_UNBLOCK, &new, NULL);
                break;
            case '1': /* Solaris */
                thr_sigsetmask(SIG_UNBLOCK, &new, NULL);
                break;
        }
        printf("\n Press CTRL-C to deliver SIGINT signal to the process\n");
        sleep(8); /* give user time to hit CTRL-C */
    }

    void
    interrupt(int sig)
    {
        printf("thread %d caught signal %d\n", thr_self( ), sig);
    }

    void test_argv(char argvl[ ]) {
        if(argvl == NULL) {
            printf("use 0 as arg1 to use thr_create( );\n \
                or use 1 as arg1 to use pthread_create( )\n");
            exit(NULL);
        }
    }
}
```

In the last example, the handler thread served as a signal-handler while also taking care of activity of its own (in this case, sleeping, although it could have been some other activity). A thread could be completely dedicated to signal-handling simply by waiting for the delivery of a selected signal by blocking with `sigwait(2)`. The two subroutines in the previous example, `handler()` and `interrupt()`, could have been replaced with the following routine:

```
void *
handler( )
{ int signal;
  printf("thread %d waiting for you to press the CTRL-C keys\n",
        thr_self( ));
  sigwait(&new, &signal);
  printf("thread %d has received the signal %d \n", thr_self( ), signal);
}
/*pthread_create( ) and thr_create( ) would use NULL instead of
```


EXAMPLE 1 Create a default thread that can serve as a signal catcher/handler with its own signal mask. (Continued)

```
    argv[1] for the arg passed to handler( ) */
In this routine, one thread is dedicated to catching and handling the signal specified by the set new, which allows main( ) and all of its other sub-threads, created after pthread_sigmask( ) or thr_sigsetmask( ) masked that signal, to continue uninterrupted. Any use of sigwait(2) should be such that all threads block the signals passed to sigwait(2) at all times. Only the thread that calls sigwait( ) will get the signals. The call to sigwait(2) takes two arguments.
```

For this type of background dedicated signal-handling routine, a Solaris daemon thread can be used by passing the argument THR_DAEMON to thr_create().

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

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

SEE ALSO sigaction(2), sigprocmask(2), sigwait(2), cond_wait(3C), pthread_create(3C), pthread_join(3C), pthread_self(3C), sigaddset(3C), sigemptyset(3C), sigsetops(3C), sleep(3C), attributes(5), standards(5)

NOTES It is not possible to block signals that cannot be ignored (see sigaction(2)). If using the threads library, it is not possible to block or unblock the signals SIGWAITING, SIGLWP, or SIGCANCEL. This restriction is quietly enforced by the threads library.

Using sigwait(2) in a dedicated thread allows asynchronously generated signals to be managed synchronously; however, sigwait(2) should never be used to manage synchronously generated signals.

Synchronously generated signals are exceptions that are generated by a thread and are directed at the thread causing the exception. Since sigwait() blocks waiting for signals, the blocking thread cannot receive a synchronously generated signal.

If sigprocmask(2) is used in a multi-threaded program, it will be the same as if thr_sigsetmask() or pthread_sigmask() has been called. POSIX leaves the semantics of the call to sigprocmask(2) unspecified in a multi-threaded process, so programs that care about POSIX portability should not depend on this semantic.

If a signal is delivered while a thread is waiting on a condition variable, the cond_wait(3C) function will be interrupted and the handler will be executed. The state of the lock protecting the condition variable is undefined while the thread is executing the signal handler.

Signals that are generated synchronously should not be masked. If such a signal is blocked and delivered, the receiving process is killed.

thr_stksegment(3C)

NAME	thr_stksegment – get thread stack address and size				
SYNOPSIS	<pre>cc -mt [flag...] file...[library...] #include <thread.h> #include <signal.h> int thr_stksegment(stack_t *ss);</pre>				
DESCRIPTION	<p>The <code>thr_stksegment()</code> function returns, in its <code>stack_t</code> argument, the address and size of the calling thread's stack.</p> <p>The <code>stack_t</code> structure includes the following members:</p> <pre>void *ss_sp size_t ss_size int ss_flags</pre> <p>On successful return from <code>thr_stksegment()</code>, <code>ss_sp</code> contains the high address of the caller's stack and <code>ss_size</code> contains the size of the stack in bytes. The <code>ss_flags</code> member is always 0. Note that the meaning of <code>ss_sp</code> is reversed from other uses of <code>stack_t</code> such as <code>sigaltstack(2)</code> where <code>ss_sp</code> is the low address.</p> <p>The stack information provided by <code>thr_stksegment()</code> is typically used by debuggers, garbage collectors, and similar applications. Most applications should not require such information.</p>				
RETURN VALUES	The <code>thr_stksegment()</code> function returns 0 if the thread stack address and size were successfully retrieved. Otherwise, it returns a non-zero error value.				
ERRORS	The <code>thr_stksegment()</code> function will fail if: EAGAIN The stack information for the thread is not available because the thread's initialization is not yet complete, or the thread is an internal thread.				
ATTRIBUTES	See <code>attributes(5)</code> 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>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>sigaltstack(2)</code> , <code>thr_create(3C)</code> , <code>attributes(5)</code>				

NAME	thr_suspend, thr_continue – suspend or continue thread execution				
SYNOPSIS	<pre>cc -mt [flag...] file... [library...] #include <thread.h> int thr_suspend(thread_t <i>target_thread</i>) ; int thr_continue(thread_t <i>target_thread</i>) ;</pre>				
DESCRIPTION	<p>The <code>thr_suspend()</code> function immediately suspends the execution of the thread specified by <i>target_thread</i>. On successful return from <code>thr_suspend()</code>, the suspended thread is no longer executing. Once a thread is suspended, subsequent calls to <code>thr_suspend()</code> have no effect.</p> <p>The <code>thr_continue()</code> function resumes the execution of a suspended thread. Once a suspended thread is continued, subsequent calls to <code>thr_continue()</code> have no effect.</p> <p>A suspended thread will not be awakened by any mechanism other than a call to <code>thr_continue()</code>. Signals and the effect of calls to <code>mutex_unlock(3C)</code>, <code>rw_unlock(3C)</code>, <code>sema_post(3C)</code>, <code>cond_signal(3C)</code>, and <code>cond_broadcast(3C)</code> remain pending until the execution of the thread is resumed by <code>thr_continue()</code>.</p>				
RETURN VALUES	If successful, the <code>thr_suspend()</code> and <code>thr_continue()</code> functions return 0. Otherwise, a non-zero value is returned to indicate the error.				
ERRORS	<p>The <code>thr_suspend()</code> and <code>thr_continue()</code> functions will fail if:</p> <p>ESRCH The <i>target_thread</i> cannot be found in the current process.</p>				
ATTRIBUTES	See <code>attributes(5)</code> 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>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>thr_create(3C)</code> , <code>thr_join(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>				
WARNINGS	<p>The <code>thr_suspend()</code> function is extremely difficult to use safely because it suspends the target thread with no concern for the target thread's state. The target thread could be holding locks, waiting for a lock, or waiting on a condition variable when it is unconditionally suspended. The thread will not run until <code>thr_continue()</code> is applied, regardless of any calls to <code>mutex_unlock()</code>, <code>cond_signal()</code>, or <code>cond_broadcast()</code> by other threads. Its existence on a sleep queue can interfere with the waking up of other threads that are on the same sleep queue.</p> <p>The <code>thr_suspend()</code> and <code>thr_continue()</code> functions should be avoided. Mechanisms that involve the cooperation of the targeted thread, such as mutex locks and condition variables, should be employed instead.</p>				

thr_yield(3C)

NAME thr_yield – yield to another thread

SYNOPSIS `cc -mt [flag...] file... [library...]`

```
#include <thread.h>

void thr_yield(void);
```

DESCRIPTION The `thr_yield()` function causes the current thread to yield its execution in favor of another thread with the same or greater priority.

RETURN VALUES The `thr_yield()` function returns nothing and does not set `errno`.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO [thr_setprio\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	times – get process times
SYNOPSIS	<pre> /usr/ucb/cc [flag ...] file ... #include <sys/param.h> #include <sys/types.h> #include <sys/times.h> int times (tmsp) ; register struct tms *tmsp; </pre>
DESCRIPTION	<p>The <code>times()</code> function returns time-accounting information for the current process and for the terminated child processes of the current process. All times are reported in clock ticks. The number of clock ticks per second is defined by the variable <code>CLK_TCK</code>, found in the header <code><limits.h></code>.</p> <p>A structure with the following members is returned by <code>times()</code>:</p> <pre> time_t tms_ftime; /* user time */ time_t tms_stime; /* system time */ time_t tms_cutime; /* user time, children */ time_t tms_cstime; /* system time, children */ </pre> <p>The children's times are the sum of the children's process times and their children's times.</p>
RETURN VALUES	Upon successful completion, <code>times()</code> returns 0. Otherwise, it returns -1.
SEE ALSO	<code>time(1)</code> , <code>time(2)</code> , <code>getrusage(3C)</code> , <code>wait(3C)</code>
NOTES	<p>Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-threaded applications is unsupported.</p> <p>The <code>times()</code> function has been superseded by <code>getrusage(3C)</code>.</p>

tmpfile(3C)

NAME	tmpfile – create a temporary file												
SYNOPSIS	<pre>#include <stdio.h> FILE *tmpfile(void);</pre>												
DESCRIPTION	<p>The tmpfile() function creates a temporary file and opens a corresponding stream. The file will automatically be deleted when all references to the file are closed. The file is opened as in fopen(3C) for update (w+).</p> <p>The largest value that can be represented correctly in an object of type off_t will be established as the offset maximum in the open file description.</p>												
RETURN VALUES	Upon successful completion, tmpfile() returns a pointer to the stream of the file that is created. Otherwise, it returns a null pointer and sets errno to indicate the error.												
ERRORS	<p>The tmpfile() function will fail if:</p> <table><tr><td>EINTR</td><td>A signal was caught during the execution of tmpfile().</td></tr><tr><td>EMFILE</td><td>There are OPEN_MAX file descriptors currently open in the calling process.</td></tr><tr><td>ENFILE</td><td>The maximum allowable number of files is currently open in the system.</td></tr><tr><td>ENOSPC</td><td>The directory or file system which would contain the new file cannot be expanded.</td></tr></table> <p>The tmpfile() function may fail if:</p> <table><tr><td>EMFILE</td><td>There are FOPEN_MAX streams currently open in the calling process.</td></tr><tr><td>ENOMEM</td><td>Insufficient storage space is available.</td></tr></table>	EINTR	A signal was caught during the execution of tmpfile().	EMFILE	There are OPEN_MAX file descriptors currently open in the calling process.	ENFILE	The maximum allowable number of files is currently open in the system.	ENOSPC	The directory or file system which would contain the new file cannot be expanded.	EMFILE	There are FOPEN_MAX streams currently open in the calling process.	ENOMEM	Insufficient storage space is available.
EINTR	A signal was caught during the execution of tmpfile().												
EMFILE	There are OPEN_MAX file descriptors currently open in the calling process.												
ENFILE	The maximum allowable number of files is currently open in the system.												
ENOSPC	The directory or file system which would contain the new file cannot be expanded.												
EMFILE	There are FOPEN_MAX streams currently open in the calling process.												
ENOMEM	Insufficient storage space is available.												
USAGE	<p>The stream refers to a file which is unlinked. If the process is killed in the period between file creation and unlinking, a permanent file may be left behind.</p> <p>The tmpfile() function has a transitional interface for 64-bit file offsets. See lf64(5).</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>Interface Stability</td><td>Standard</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard								
ATTRIBUTE TYPE	ATTRIBUTE VALUE												
Interface Stability	Standard												
SEE ALSO	unlink(2), fopen(3C), mkstemp(3C), mktemp(3C), tmpnam(3C), lf64(5), standards(5)												

NAME	tmpnam, tmpnam_r, tmpnam – create a name for a temporary file
SYNOPSIS	<pre>#include <stdio.h> char *tmpnam(char *s); char *tmpnam_r(char *s); char *tempnam(const char *dir, const char *pfx);</pre>
DESCRIPTION	<p>These functions generate file names that can be used safely for a temporary file.</p> <p>tmpnam() The tmpnam() function always generates a file name using the path prefix defined as P_tmpdir in the <stdio.h> header. On Solaris systems, the default value for P_tmpdir is /var/tmp. If s is NULL, tmpnam() leaves its result in a thread-specific data area and returns a pointer to that area. The next call to tmpnam() by the same thread will destroy the contents of the area. If s is not NULL, it is assumed to be the address of an array of at least L_tmpnam bytes, where L_tmpnam is a constant defined through inclusion of <stdio.h>. The tmpnam() function places its result in that array and returns s.</p> <p>tmpnam_r() The tmpnam_r() function has the same functionality as tmpnam() except that if s is a null pointer, the function returns NULL.</p> <p>tempnam() The tempnam() function allows the user to control the choice of a directory. The argument dir points to the name of the directory in which the file is to be created. If dir is NULL or points to a string that is not a name for an appropriate directory, the path prefix defined as P_tmpdir in the <stdio.h> header is used. If that directory is not accessible, /tmp is used. If, however, the TMPDIR environment variable is set in the user's environment, its value is used as the temporary-file directory.</p> <p>Many applications prefer that temporary files have certain initial character sequences in their names. The pfx argument may be NULL or point to a string of up to five characters to be used as the initial characters of the temporary-file name.</p> <p>Upon successful completion, tempnam() uses malloc(3C) to allocate space for a string, puts the generated pathname in that space, and returns a pointer to it. The pointer is suitable for use in a subsequent call to free(). If tempnam() cannot return the expected result for any reason (for example, malloc() failed), or if none of the above-mentioned attempts to find an appropriate directory was successful, a null pointer is returned and errno is set to indicate the error.</p>
ERRORS	<p>The tempnam() function will fail if:</p> <p>ENOMEM Insufficient storage space is available.</p>
USAGE	<p>These functions generate a different file name each time they are called.</p> <p>Files created using these functions and either fopen(3C) or creat(2) are temporary only in the sense that they reside in a directory intended for temporary use, and their names are unique. It is the user's responsibility to remove the file when its use is ended.</p>

tmpnam(3C)

If called more than `TMP_MAX` (defined in `<stdio.h>`) times in a single process, these functions start recycling previously used names.

Between the time a file name is created and the file is opened, it is possible for some other process to create a file with the same name. This can never happen if that other process is using these functions or [mktemp\(3C\)](#) and the file names are chosen to render duplication by other means unlikely.

The `tmpnam()` function is safe to use in multithreaded applications because it employs thread-specific data if it is passed a NULL pointer. However, its use is discouraged. The `tempnam()` function is safe in multithreaded applications and should be used instead.

When compiling multithreaded applications, the `_REENTRANT` flag must be defined on the compile line. This flag should be used only with multithreaded applications.

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	<code>tmpnam()</code> and <code>tempnam()</code> are Standard.
MT-Level	Safe

SEE ALSO [creat\(2\)](#), [unlink\(2\)](#), [fopen\(3C\)](#), [free\(3C\)](#), [malloc\(3C\)](#), [mktemp\(3C\)](#), [mkstemp\(3C\)](#), [tmpfile\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME toascii – translate integer to a 7-bit ASCII character

SYNOPSIS `#include <ctype.h>`
`int toascii(int c);`

DESCRIPTION The `toascii()` function converts its argument into a 7-bit ASCII character.

RETURN VALUES The `toascii()` function returns the value $(c \ \& \ 0x7f)$.

ERRORS No errors are returned.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `isascii(3C)`, `attributes(5)`, `standards(5)`

`_tolower(3C)`

NAME	<code>_tolower</code> – transliterate upper-case characters to lower-case
SYNOPSIS	<pre>#include <ctype.h> int _tolower(int c);</pre>
DESCRIPTION	The <code>_tolower()</code> macro is equivalent to <code>tolower(3C)</code> except that the argument <code>c</code> must be an upper-case letter.
RETURN VALUES	On successful completion, <code>_tolower()</code> returns the lower-case letter corresponding to the argument passed.
ERRORS	No errors are defined.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [isupper\(3C\)](#), [tolower\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	tolower – transliterate upper-case characters to lower-case
SYNOPSIS	<pre>#include <ctype.h> int tolower(int c);</pre>
DESCRIPTION	The <code>tolower()</code> function has as a domain a type <code>int</code> , the value of which is representable as an unsigned <code>char</code> or the value of <code>EOF</code> . If the argument has any other value, the argument is returned unchanged. If the argument of <code>tolower()</code> represents an upper-case letter, and there exists a corresponding lower-case letter (as defined by character type information in the program locale category <code>LC_CTYPE</code>), the result is the corresponding lower-case letter. All other arguments in the domain are returned unchanged.
RETURN VALUES	On successful completion, <code>tolower()</code> returns the lower-case letter corresponding to the argument passed. Otherwise, it returns the argument unchanged.
ERRORS	No errors are defined.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [_tolower\(3C\)](#), [setlocale\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

`_toupper(3C)`

NAME	<code>_toupper</code> – transliterate lower-case characters to upper-case
SYNOPSIS	<pre>#include <ctype.h> int _toupper(int c);</pre>
DESCRIPTION	The <code>_toupper()</code> macro is equivalent to <code>toupper(3C)</code> except that the argument <code>c</code> must be a lower-case letter.
RETURN VALUES	On successful completion, <code>_toupper()</code> returns the upper-case letter corresponding to the argument passed.
ERRORS	No errors are defined.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `islower(3C)`, `toupper(3C)`, `attributes(5)`, `standards(5)`

NAME	toupper – transliterate lower-case characters to upper-case
SYNOPSIS	<pre>#include <ctype.h> int toupper(int c);</pre>
DESCRIPTION	The <code>toupper()</code> function has as a domain a type <code>int</code> , the value of which is representable as an unsigned <code>char</code> or the value of <code>EOF</code> . If the argument has any other value, the argument is returned unchanged. If the argument of <code>toupper()</code> represents a lower-case letter, and there exists a corresponding upper-case letter (as defined by character type information in the program locale category <code>LC_CTYPE</code>), the result is the corresponding upper-case letter. All other arguments in the domain are returned unchanged.
RETURN VALUES	On successful completion, <code>toupper()</code> returns the upper-case letter corresponding to the argument passed.
ERRORS	No errors are defined.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [_toupper\(3C\)](#), [setlocale\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

towctrans(3C)

NAME	towctrans – wide-character mapping								
SYNOPSIS	<pre>#include <wctype.h> wint_t towctrans(wint_t <i>wc</i>, wctrans_t <i>desc</i>);</pre>								
DESCRIPTION	<p>The <code>towctrans()</code> function maps the wide character <code>wc</code> using the mapping described by <code>desc</code>. The current setting of the <code>LC_CTYPE</code> category shall be the same as during the call to <code>wctrans()</code> that returned the value <code>desc</code>.</p> <p>The function call <code>towctrans(wc, wctrans("tolower"))</code> behaves the same as <code>tolower(wc)</code>.</p> <p>The function call <code>towctrans(wc, wctrans("toupper"))</code> behaves the same as <code>toupper(wc)</code>.</p>								
RETURN VALUES	The <code>towctrans()</code> function returns the mapped value of <code>wc</code> , using the mapping described by <code>desc</code> ; otherwise, it returns <code>wc</code> unchanged.								
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:								
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>CSI</td><td>Enabled</td></tr><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe with exceptions</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	Standard	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
CSI	Enabled								
Interface Stability	Standard								
MT-Level	MT-Safe with exceptions								
SEE ALSO	<code>setlocale(3C)</code> , <code>wctrans(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>								

NAME	towlower – transliterate upper-case wide-character code to lower-case
SYNOPSIS	<pre>#include <wchar.h> wint_t tolower(wint_t wc);</pre>
DESCRIPTION	The <code>towlower()</code> function has as a domain a type <code>wint_t</code> , the value of which must be a character representable as a <code>wchar_t</code> , and must be a wide-character code corresponding to a valid character in the current locale or the value of <code>WEOF</code> . If the argument has any other value, the argument is returned unchanged. If the argument of <code>towlower()</code> represents an upper-case wide-character code, and there exists a corresponding lower-case wide-character code (as defined by character type information in the program locale category <code>LC_CTYPE</code>), the result is the corresponding lower-case wide-character code. All other arguments in the domain are returned unchanged.
RETURN VALUES	On successful completion, <code>towlower()</code> returns the lower-case letter corresponding to the argument passed. Otherwise, it returns the argument unchanged.
ERRORS	No errors are defined.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [iswalph\(3C\)](#), [setlocale\(3C\)](#), [towupper\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

towupper(3C)

NAME	towupper – transliterate lower-case wide-character code to upper-case								
SYNOPSIS	<pre>#include <wchar.h> wint_t towupper(wint_t wc);</pre>								
DESCRIPTION	The <code>towupper()</code> function has as a domain a type <code>wint_t</code> , the value of which must be a character representable as a <code>wchar_t</code> , and must be a wide-character code corresponding to a valid character in the current locale or the value of <code>WEOF</code> . If the argument has any other value, the argument is returned unchanged. If the argument of <code>towupper()</code> represents a lower-case wide-character code (as defined by character type information in the program locale category <code>LC_CTYPE</code>), the result is the corresponding upper-case wide-character code. All other arguments in the domain are returned unchanged.								
RETURN VALUES	Upon successful completion, <code>towupper()</code> returns the upper-case letter corresponding to the argument passed. Otherwise, it returns the argument unchanged.								
ERRORS	No errors are defined.								
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:								
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>CSI</td><td>Enabled</td></tr><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
CSI	Enabled								
Interface Stability	Standard								
MT-Level	MT-Safe								
SEE ALSO	<code>iswalpha(3C)</code> , <code>setlocale(3C)</code> , <code>towlower(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>								

NAME	truncate, ftruncate – set a file to a specified length
SYNOPSIS	<pre>#include <unistd.h> int truncate(const char *path, off_t length); int ftruncate(int fildes, off_t length);</pre>
DESCRIPTION	<p>The <code>truncate()</code> function causes the regular file named by <i>path</i> to have a size equal to <i>length</i> bytes.</p> <p>If the file previously was larger than <i>length</i>, the extra data is discarded. If the file was previously shorter than <i>length</i>, its size is increased, and the extended area appears as if it were zero-filled.</p> <p>The application must ensure that the process has write permission for the file.</p> <p>This function does not modify the file offset for any open file descriptions associated with the file.</p> <p>The <code>ftruncate()</code> function causes the regular file referenced by <i>fildes</i> to be truncated to <i>length</i>. If the size of the file previously exceeded <i>length</i>, the extra data is no longer available to reads on the file. If the file previously was smaller than this size, <code>ftruncate()</code> increases the size of the file with the extended area appearing as if it were zero-filled. The value of the seek pointer is not modified by a call to <code>ftruncate()</code>.</p> <p>The <code>ftruncate()</code> function works only with regular files and shared memory. If <i>fildes</i> refers to a shared memory object, <code>ftruncate()</code> sets the size of the shared memory object to <i>length</i>. If <i>fildes</i> refers to a directory or is not a valid file descriptor open for writing, <code>ftruncate()</code> fails.</p> <p>If the effect of <code>ftruncate()</code> is to decrease the size of a shared memory object or memory mapped file and whole pages beyond the new end were previously mapped, then the whole pages beyond the new end shall be discarded.</p> <p>If the effect of <code>ftruncate()</code> is to increase the size of a shared memory object, it is unspecified if the contents of any mapped pages between the old end-of-file and the new are flushed to the underlying object.</p> <p>These functions do not modify the file offset for any open file descriptions associated with the file. On successful completion, if the file size is changed, these functions will mark for update the <code>st_ctime</code> and <code>st_mtime</code> fields of the file, and if the file is a regular file, the <code>S_ISUID</code> and <code>S_ISGID</code> bits of the file mode are left unchanged.</p> <p>If the request would cause the file size to exceed the soft file size limit for the process, the request will fail and a <code>SIGXFSZ</code> signal will be generated for the process.</p>
RETURN VALUES	Upon successful completion, <code>ftruncate()</code> and <code>truncate()</code> return 0. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.
ERRORS	The <code>ftruncate()</code> and <code>truncate()</code> functions will fail if:

truncate(3C)

EINTR	A signal was caught during execution.
EINVAL	The <i>length</i> argument was less than 0.
EFBIG or EINVAL	The <i>length</i> argument was greater than the maximum file size.
EIO	An I/O error occurred while reading from or writing to a file system.
EROFS	The named file resides on a read-only file system.
The <code>truncate()</code> function will fail if:	
EACCES	A component of the path prefix denies search permission, or write permission is denied on the file.
EFAULT	The <i>path</i> argument points outside the process' allocated address space.
EINVAL	The <i>path</i> argument is not an ordinary file.
EISDIR	The named file is a directory.
ELOOP	Too many symbolic links were encountered in resolving <i>path</i> .
EMFILE	The maximum number of file descriptors available to the process has been reached.
ENAMETOOLONG	The length of the specified pathname exceeds {PATH_MAX} bytes, or the length of a component of the pathname exceeds {NAME_MAX} bytes.
ENOENT	A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.
ENFILE	Additional space could not be allocated for the system file table.
ENOTDIR	A component of the path prefix of <i>path</i> is not a directory.
ENOLINK	The <i>path</i> argument points to a remote machine and the link to that machine is no longer active.
The <code>ftruncate()</code> function will fail if:	
EAGAIN	The file exists, mandatory file/record locking is set, and there are outstanding record locks on the file (see <code>chmod(2)</code>).
EBADF or EINVAL	The <i>fil-des</i> argument is not a file descriptor open for writing.
EFBIG	The file is a regular file and <i>length</i> is greater than the offset maximum established in the open file description associated with <i>fil-des</i> .
EINVAL	The <i>fil-des</i> argument references a file that was opened without write permission.
EINVAL	The <i>fil-des</i> argument does not correspond to an ordinary file.

ENOLINK The *fildev* argument points to a remote machine and the link to that machine is no longer active.

The `truncate()` function may fail if:

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

USAGE The `truncate()` and `ftruncate()` 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
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `chmod(2)`, `fcntl(2)`, `open(2)`, `attributes(5)`, `lf64(5)`, `standards(5)`

tsearch(3C)

NAME	tsearch, tfind, tdelete, twalk – manage binary search trees
SYNOPSIS	<pre>#include <search.h> void *tsearch(const void *key, void **rootp, int (*compar) (const void *, const void *)); void *tfind(const void *key, void * const *rootp, int (*compar) (const void *, const void *)); void *tdelete(const void *restrict key, void **restrict rootp, int (*compar) (const void *, const void *)); void twalk(const void *root, void(*action) (void *, VISIT, int));</pre>
DESCRIPTION	<p>The <code>tsearch()</code>, <code>tfind()</code>, <code>tdelete()</code>, and <code>twalk()</code> functions are routines for manipulating binary search trees. They are generalized from <i>Knuth (6.2.2) Algorithms T and D</i>. All comparisons are done with a user-supplied routine. This routine is called with two arguments, the pointers to the elements being compared. It returns an integer less than, equal to, or greater than 0, according to whether the first argument is to be considered less than, equal to or greater than the second argument. The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.</p> <p>The <code>tsearch()</code> function is used to build and access the tree. The <code>key</code> argument is a pointer to a datum to be accessed or stored. If there is a datum in the tree equal to <code>*key</code> (the value pointed to by <code>key</code>), a pointer to this found datum is returned. Otherwise, <code>*key</code> is inserted, and a pointer to it returned. Only pointers are copied, so the calling routine must store the data. The <code>rootp</code> argument points to a variable that points to the root of the tree. A null value for the variable pointed to by <code>rootp</code> denotes an empty tree; in this case, the variable will be set to point to the datum which will be at the root of the new tree.</p> <p>Like <code>tsearch()</code>, <code>tfind()</code> will search for a datum in the tree, returning a pointer to it if found. However, if it is not found, <code>tfind()</code> will return a null pointer. The arguments for <code>tfind()</code> are the same as for <code>tsearch()</code>.</p> <p>The <code>tdelete()</code> function deletes a node from a binary search tree. The arguments are the same as for <code>tsearch()</code>. The variable pointed to by <code>rootp</code> will be changed if the deleted node was the root of the tree. <code>tdelete()</code> returns a pointer to the parent of the deleted node, or a null pointer if the node is not found.</p> <p>The <code>twalk()</code> function traverses a binary search tree. The <code>root</code> argument is the root of the tree to be traversed. (Any node in a tree may be used as the root for a walk below that node.) <code>action</code> is the name of a routine to be invoked at each node. This routine is, in turn, called with three arguments. The first argument is the address of the node being visited. The second argument is a value from an enumeration data type</p> <pre>typedef enum { preorder, postorder, endorder, leaf } VISIT;</pre>

(defined in `<search.h>`), depending on whether this is the first, second or third time that the node has been visited (during a depth-first, left-to-right traversal of the tree), or whether the node is a leaf. The third argument is the level of the node in the tree, with the root being level zero.

The pointers to the key and the root of the tree should be of type pointer-to-element, and cast to type pointer-to-character. Similarly, although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

RETURN VALUES

If the node is found, both `tsearch()` and `tfind()` return a pointer to it. If not, `tfind()` returns a null pointer, and `tsearch()` returns a pointer to the inserted item.

A null pointer is returned by `tsearch()` if there is not enough space available to create a new node.

A null pointer is returned by `tsearch()`, `tfind()` and `tdelete()` if `rootp` is a null pointer on entry.

The `tdelete()` function returns a pointer to the parent of the deleted node, or a null pointer if the node is not found.

The `twalk()` function returns no value.

ERRORS

No errors are defined.

USAGE

The `root` argument to `twalk()` is one level of indirection less than the `rootp` arguments to `tsearch()` and `tdelete()`.

There are two nomenclatures used to refer to the order in which tree nodes are visited. `tsearch()` uses preorder, postorder and endorder to refer respectively to visiting a node before any of its children, after its left child and before its right, and after both its children. The alternate nomenclature uses preorder, inorder and postorder to refer to the same visits, which could result in some confusion over the meaning of postorder.

If the calling function alters the pointer to the root, the results are unpredictable.

These functions safely allows concurrent access by multiple threads to disjoint data, such as overlapping subtrees or tables.

EXAMPLES

EXAMPLE 1 A sample program of using `tsearch()` function.

The following code reads in strings and stores structures containing a pointer to each string and a count of its length. It then walks the tree, printing out the stored strings and their lengths in alphabetical order.

```
#include <string.h>
#include <stdio.h>
#include <search.h>
struct node {
    char *string;
    int length;
```

tsearch(3C)

EXAMPLE 1 A sample program of using `tsearch()` function. *(Continued)*

```
};
char string_space[10000];
struct node nodes[500];
void *root = NULL;

int node_compare(const void *node1, const void *node2) {
    return strcmp(((const struct node *) node1)->string,
                 ((const struct node *) node2)->string);
}

void print_node(const void *node, VISIT order, int level) {
    if (order == preorder || order == leaf) {
        printf("length=%d, string=%20s\n",
              (*(struct node **)node)->length,
              (*(struct node **)node)->string);
    }
}

main()
{
    char *strptr = string_space;
    struct node *nodeptr = nodes;
    int i = 0;

    while (gets(strptr) != NULL && i++ < 500) {
        nodeptr->string = strptr;
        nodeptr->length = strlen(strptr);
        (void) tsearch((void *)nodeptr,
                      &root, node_compare);
        strptr += nodeptr->length + 1;
        nodeptr++;
    }
    twalk(root, print_node);
}
```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [bsearch\(3C\)](#), [hsearch\(3C\)](#), [lsearch\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME | ttynam, ttynam_r – find pathname of a terminal

SYNOPSIS | #include <unistd.h>
char *ttynam(int *filde*);
char *ttynam_r(int *filde*, char **name*, int *namelen*);

Standard conforming | cc [*flag ...*] *file ...* -D_POSIX_PTHREAD_SEMANTICS [*library ...*]
int ttynam_r(int *filde*, char **name*, size_t *namesize*);

DESCRIPTION | The ttynam() function returns a pointer to a string containing the null-terminated path name of the terminal device associated with file descriptor *filde*. The return value points to thread-specific data whose content is overwritten by each call from the same thread.
The ttynam_r() function has the same functionality as ttynam() except that the caller must supply a buffer *name* with length *namelen* to store the result; this buffer must be at least _POSIX_PATH_MAX in size (defined in <limits.h>). The standard-conforming version (see standards(5)) of ttynam_r() takes a *namesize* parameter of type size_t.

RETURN VALUES | Upon successful completion, ttynam() and ttynam_r() return a pointer to a string. Otherwise, a null pointer is returned and *errno* is set to indicate the error.
The standard-conforming ttynam_r() returns 0 if successful or the error number upon failure.

ERRORS | The ttynam() and ttynam_r() functions may fail if:
EBADF The *filde* argument is not a valid file descriptor. This condition is reported.
ENOTTY The *filde* argument does not refer to a terminal device. This condition is reported.
The ttynam_r() function may fail if:
ERANGE The value of *namesize* is smaller than the length of the string to be returned including the terminating null character.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Safe

SEE ALSO | Intro(3), gettext(3C), setlocale(3C), attributes(5), standards(5)

NOTES | When compiling multithreaded applications, see Intro(3), Notes On Multithreaded Applications.

ttyname(3C)

If the application is linked with `-lintl`, then messages printed from this function are in the native language specified by the `LC_MESSAGES` locale category; see [setlocale\(3C\)](#).

The return value of `ttyname()` points to thread-specific data whose content is overwritten by each call from the same thread. This function is safe to use in multithreaded applications, but its use is discouraged. The `ttyname_r()` function should be used instead.

Solaris 2.4 and earlier releases provided definitions of the `ttyname_r()` interface 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 might not be supported in future releases. New applications and libraries should use the standard-conforming interface.

NAME	ttyslot – find the slot of the current user in the user accounting database				
SYNOPSIS	<pre>#include <stdlib.h> int ttyslot(void);</pre>				
DESCRIPTION	The <code>ttyslot()</code> function returns the index of the current user's entry in the user accounting database, <code>/var/adm/utmpx</code> . The current user's entry is an entry for which the <code>utline</code> member matches the name of a terminal device associated with any of the process's file descriptors 0, 1 or 2. The index is an ordinal number representing the record number in the database of the current user's entry. The first entry in the database is represented by the return value 0.				
RETURN VALUES	Upon successful completion, <code>ttyslot()</code> returns the index of the current user's entry in the user accounting database. If an error was encountered while searching for the terminal name or if none of the above file descriptors are associated with a terminal device, <code>-1</code> is returned.				
FILES	<code>/var/adm/utmpx</code> user access and accounting information				
ATTRIBUTES	See <code>attributes(5)</code> 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>Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	Safe				
SEE ALSO	<code>getutent(3C)</code> , <code>ttyname(3C)</code> , <code>utmpx(4)</code> , <code>attributes(5)</code>				

ualarm(3C)

NAME	ualarm – schedule signal after interval in microseconds				
SYNOPSIS	<pre>#include <unistd.h> useconds_t ualarm(useconds_t <i>useconds</i>, useconds_t <i>interval</i>);</pre>				
DESCRIPTION	<p>The <code>ualarm()</code> function causes the <code>SIGALRM</code> signal to be generated for the calling process after the number of real-time microseconds specified by the <i>useconds</i> argument has elapsed. When the <i>interval</i> argument is non-zero, repeated timeout notification occurs with a period in microseconds specified by the <i>interval</i> argument. If the notification signal, <code>SIGALRM</code>, is not caught or ignored, the calling process is terminated.</p> <p>Because of scheduling delays, resumption of execution when the signal is caught may be delayed an arbitrary amount of time.</p> <p>Interactions between <code>ualarm()</code> and either <code>alarm(2)</code> or <code>sleep(3C)</code> are unspecified.</p>				
RETURN VALUES	The <code>ualarm()</code> function returns the number of microseconds remaining from the previous <code>ualarm()</code> call. If no timeouts are pending or if <code>ualarm()</code> has not previously been called, <code>ualarm()</code> returns 0.				
ERRORS	No errors are defined.				
USAGE	The <code>ualarm()</code> function is a simplified interface to <code>setitimer(2)</code> , and uses the <code>ITIMER_REAL</code> interval timer.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Interface Stability	Standard				
SEE ALSO	<code>alarm(2)</code> , <code>setitimer(2)</code> , <code>sighold(3C)</code> , <code>signal(3C)</code> , <code>sleep(3C)</code> , <code>usleep(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>				

NAME	ucred_get, ucred_free, ucred_geteuid, ucred_getruid, ucred_getsuid, ucred_getegid, ucred_getrgid, ucred_getsgid, ucred_getgroups, ucred_getprivset, ucred_getpid, ucred_getprojid, ucred_getzoneid, ucred_getpflags, ucred_size – user credential functions
SYNOPSIS	<pre>#include <ucred.h> ucred_t *ucred_get(pid_t pid); void ucred_free(const ucred_t *uc); uid_t ucred_geteuid(const ucred_t *uc); uid_t ucred_getruid(const ucred_t *uc); uid_t ucred_getsuid(const ucred_t *uc); gid_t ucred_getegid(const ucred_t *uc); gid_t ucred_getrgid(const ucred_t *uc); gid_t ucred_getsgid(const ucred_t *uc); int ucred_getgroups(const ucred_t *uc, const gid_t **groups); const priv_set_t *ucred_getprivset(const ucred_t *uc, const char * set); pid_t ucred_getpid(const ucred_t *uc); projid_t ucred_getprojid(const ucred_t *uc); zoneid_t ucred_getzoneid(const ucred_t *uc); uint_t ucred_getpflags(const ucred_t *uc, uint_t flags); size_t ucred_size(void);</pre>
DESCRIPTION	<p>These functions return or act on a user credential, <code>ucred_t</code>. User credentials are returned by various functions and describe the credentials of a process. Information about the process can then be obtained by calling the access functions. Access functions can fail if the underlying mechanism did not return sufficient information.</p> <p>The <code>ucred_get()</code> function returns the user credential of the specified <code>pid</code> or <code>NULL</code> if none can be obtained. A <code>pid</code> value of <code>P_MYID</code> returns information about the calling process. The return value is dynamically allocated and must be freed using <code>ucred_free()</code>.</p> <p>The <code>ucred_geteuid()</code>, <code>ucred_getruid()</code>, <code>ucred_getsuid()</code>, <code>ucred_getegid()</code>, <code>ucred_getrgid()</code>, and <code>ucred_getsgid()</code> functions return the effective UID, real UID, saved UID, effective GID, real GID, saved GID, respectively, or -1 if the user credential does not contain sufficient information.</p> <p>The <code>ucred_getgroups()</code> function stores a pointer to the group list in the <code>gid_t *</code> pointed to by the second argument and returns the number of groups in the list. It returns -1 if the information is not available. The returned group list is valid until <code>ucred_free()</code> is called on the user credential given as argument.</p>

ucred_get(3C)

The `ucred_getpid()` function returns the process ID of the process or -1 if the process ID is not available. The process ID returned in a user credential is only guaranteed to be correct in a very limited number of cases when returned by `door_ucred(3DOOR)` and `ucred_get()`. In all other cases, the process in question might have handed off the file descriptor, the process might have exited or executed another program, or the process ID might have been reused by a completely unrelated process after the original program exited.

The `ucred_getprojid()` function returns the project ID of the process or -1 if the project ID is not available.

The `ucred_getzoneid()` function returns the zone ID of the process or -1 if the zone ID is not available.

The `ucred_getprivset()` function returns the specified privilege set specified as second argument, or NULL if either the requested information is not available or the `priv_ptype_t` argument is invalid. The returned privilege set is valid until `ucred_free()` is called on the specified user credential.

The `ucred_getpflags()` function returns the value of the specified privilege flags from the `ucred` structure, or `(uint_t)-1` if none was present.

The `ucred_free()` function frees the memory allocated for the specified user credential.

The `ucred_size()` function returns `sizeof(ucred_t)`. This value is constant only until the next boot, at which time it could change. The `ucred_size()` function can be used to determine the size of the buffer needed to receive a credential option with `SO_RECVUCRED`. See `socket.h(3HEAD)`.

RETURN VALUES

See DESCRIPTION.

ERRORS

The `ucred_get()` function will fail if:

EAGAIN	There is not enough memory available to allocate sufficient memory to hold a user credential. The application can try again later.
EACCES	The caller does not have sufficient privileges to examine the target process.
EMFILE	The calling process cannot open any more files.
ENFILE	
ENOMEM	The physical limits of the system are exceeded by the memory allocation needed to hold a user credential.
ESRCH	The target process does not exist.

The `ucred_getprivset()` function will fail if:

EINVAL	The privilege set argument is invalid.
--------	--

ucred_get(3C)

The `ucred_geteuid()`, `ucred_getruid()`, `ucred_getsuid()`, `ucred_getegid()`, `ucred_getrgid()`, `ucred_getsgid()`, `ucred_getgroups()`, `ucred_getpflags()`, `ucred_getprivset()`, `ucred_getprojid()`, and `ucred_getpid()` functions will fail if:

EINVAL The requested user credential attribute is not available in the specified user credential.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving
MT-Level	MT-Safe

SEE ALSO `getpflags(2)`, `getppriv(2)`, `door_ucred(3DOOR)`, `priv_set(3C)`, `socket.h(3HEAD)`, `attributes(5)`, `privileges(5)`

umem_alloc(3MALLOC)

NAME	umem_alloc, umem_zalloc, umem_free, umem_nofail_callback – fast, scalable memory allocation				
SYNOPSIS	<pre>cc [<i>flag</i> ...] <i>file</i>... -lumem [<i>library</i> ...] #include <umem.h> void *umem_alloc(size_t size, int flags); void *umem_zalloc(size_t size, int flags); void umem_free(void *buf, size_t size); void umem_nofail_callback((int (*callback)(void)); void *malloc(size_t size); void *calloc(size_t nelem, size_t elsize); void free(void *ptr); void *memalign(size_t alignment, size_t size); void *realloc(void *ptr, size_t size); void *valloc(size_t size);</pre>				
DESCRIPTION	<p>The <code>umem_alloc()</code> function returns a pointer to a block of <i>size</i> bytes suitably aligned for any variable type. The initial contents of memory allocated using <code>umem_alloc()</code> is undefined. The <i>flags</i> argument determines the behavior of <code>umem_alloc()</code> if it is unable to fulfill the request. The <i>flags</i> argument can take the following values:</p> <table><tr><td>UMEM_DEFAULT</td><td>Return NULL on failure.</td></tr><tr><td>UMEM_NOFAIL</td><td>Call an optional <i>callback</i> (set with <code>umem_nofail_callback()</code>) on failure. The <i>callback</i> takes no arguments and can finish by:<ul style="list-style-type: none">■ returning <code>UMEM_CALLBACK_RETRY</code>, in which case the allocation will be retried. If the allocation fails, the callback will be invoked again.■ returning <code>UMEM_CALLBACK_EXIT(status)</code>, in which case <code>exit(2)</code> is invoked with <i>status</i> as its argument. The <code>exit()</code> function is called only once. If multiple threads return from the <code>UMEM_NOFAIL</code> callback with <code>UMEM_CALLBACK_EXIT(status)</code>, one will call <code>exit()</code> while the other blocks until <code>exit()</code> terminates the program.■ invoking a context-changing function (<code>setcontext(2)</code>) or a non-local jump (<code>longjmp(3C)</code> or <code>siglongjmp(3C)</code>), or ending the current thread of control (<code>thr_exit(3C)</code> or <code>pthread_exit(3C)</code>). The application is responsible for any necessary cleanup. The state of <code>libumem</code> remains consistent.</td></tr></table> <p>If no callback has been set or the callback has been set to NULL, <code>umem_alloc(..., UMEM_NOFAIL)</code> behaves as though the callback returned <code>UMEM_CALLBACK_EXIT(255)</code>.</p>	UMEM_DEFAULT	Return NULL on failure.	UMEM_NOFAIL	Call an optional <i>callback</i> (set with <code>umem_nofail_callback()</code>) on failure. The <i>callback</i> takes no arguments and can finish by: <ul style="list-style-type: none">■ returning <code>UMEM_CALLBACK_RETRY</code>, in which case the allocation will be retried. If the allocation fails, the callback will be invoked again.■ returning <code>UMEM_CALLBACK_EXIT(status)</code>, in which case <code>exit(2)</code> is invoked with <i>status</i> as its argument. The <code>exit()</code> function is called only once. If multiple threads return from the <code>UMEM_NOFAIL</code> callback with <code>UMEM_CALLBACK_EXIT(status)</code>, one will call <code>exit()</code> while the other blocks until <code>exit()</code> terminates the program.■ invoking a context-changing function (<code>setcontext(2)</code>) or a non-local jump (<code>longjmp(3C)</code> or <code>siglongjmp(3C)</code>), or ending the current thread of control (<code>thr_exit(3C)</code> or <code>pthread_exit(3C)</code>). The application is responsible for any necessary cleanup. The state of <code>libumem</code> remains consistent.
UMEM_DEFAULT	Return NULL on failure.				
UMEM_NOFAIL	Call an optional <i>callback</i> (set with <code>umem_nofail_callback()</code>) on failure. The <i>callback</i> takes no arguments and can finish by: <ul style="list-style-type: none">■ returning <code>UMEM_CALLBACK_RETRY</code>, in which case the allocation will be retried. If the allocation fails, the callback will be invoked again.■ returning <code>UMEM_CALLBACK_EXIT(status)</code>, in which case <code>exit(2)</code> is invoked with <i>status</i> as its argument. The <code>exit()</code> function is called only once. If multiple threads return from the <code>UMEM_NOFAIL</code> callback with <code>UMEM_CALLBACK_EXIT(status)</code>, one will call <code>exit()</code> while the other blocks until <code>exit()</code> terminates the program.■ invoking a context-changing function (<code>setcontext(2)</code>) or a non-local jump (<code>longjmp(3C)</code> or <code>siglongjmp(3C)</code>), or ending the current thread of control (<code>thr_exit(3C)</code> or <code>pthread_exit(3C)</code>). The application is responsible for any necessary cleanup. The state of <code>libumem</code> remains consistent.				

umem_alloc(3MALLOC)

The libumem library can call callbacks from any place that a UMEM_NOFAIL allocation is issued. In multithreaded applications, callbacks are expected to perform their own concurrency management.

The function call `umem_alloc(0, flag)` always returns `NULL`. The function call `umem_free(NULL, 0)` is allowed.

The `umem_zalloc()` function has the same semantics as `umem_alloc()`, but the block of memory is initialized to zeros before it is returned.

The `umem_free()` function frees blocks previously allocated using `umem_alloc()` and `umem_zalloc()`. The buffer address and size must exactly match the original allocation. Memory must not be returned piecemeal.

The `umem_nofail_callback()` function sets the process-wide UMEM_NOFAIL callback. See the description of UMEM_NOFAIL for more information.

The `malloc()`, `calloc()`, `free()`, `memalign()`, `realloc()`, and `valloc()` functions are as described in [malloc\(3C\)](#). The libumem library provides these functions for backwards-compatibility with the standard functions.

ENVIRONMENT VARIABLES

See [umem_debug\(3MALLOC\)](#) for environment variables that effect the debugging features of the libumem library.

UMEM_OPTIONS Contains a list of comma-separated options. Unrecognized options are ignored. The options that are supported are:

`backend=sbrk`

`backend=mmap`

Set the underlying function used to allocate memory. This option can be set to `sbrk` (the default) for an `sbrk(2)`-based source or `mmap` for an `mmap(2)`-based source. If set to a value that is not supported, `sbrk` will be used.

EXAMPLES

EXAMPLE 1 Using the `umem_alloc()` function.

```
#include <stdio.h>
#include <umem.h>
...
char *buf = umem_alloc(1024, UMEM_DEFAULT);

if (buf == NULL) {
    fprintf(stderr, "out of memory\n");
    return (1);
}
/* cannot assume anything about buf's contents */
...
umem_free(buf, 1024);
...
```

umem_alloc(3MALLOC)

EXAMPLE 2 Using the umem_zalloc() function

```
#include <stdio.h>
#include <umem.h>
...
char *buf = umem_zalloc(1024, UMEM_DEFAULT);

if (buf == NULL) {
    fprintf(stderr, "out of memory\n");
    return (1);
}
/* buf contains zeros */
...
umem_free(buf, 1024);
...
```

EXAMPLE 3 Using UMEM_NOFAIL

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

/*
 * Note that the allocation code below does not have to
 * check for umem_alloc() returning NULL
 */
int
my_failure_handler(void)
{
    (void) fprintf(stderr, "out of memory\n");
    return (UMEM_CALLBACK_EXIT(255));
}
...
umem_nofail_callback(my_failure_handler);
...
int i;
char *buf[100];

for (i = 0; i < 100; i++)
    buf[i] = umem_alloc(1024 * 1024, UMEM_NOFAIL);
...
for (i = 0; i < 100; i++)
    umem_free(buf[i], 1024 * 1024);
...
```

EXAMPLE 4 Using UMEM_NOFAIL in a multithreaded application

```
#define _REENTRANT
#include <thread.h>
#include <stdio.h>
#include <umem.h>

void *
start_func(void *the_arg)
{
    int *info = (int *)the_arg;
```


EXAMPLE 4 Using UMEM_NOFAIL in a multithreaded application (Continued)

```

        char *buf = umem_alloc(1024 * 1024, UMEM_NOFAIL);

        /* does not need to check for buf == NULL */
        buf[0] = 0;
        ...
        /*
         * if there were other UMEM_NOFAIL allocations,
         * we would need to arrange for buf to be
         * umem_free()ed upon failure.
         */
        ...
        umem_free(buf, 1024 * 1024);
        return (the_arg);
    }
    ...
    int
    my_failure_handler(void)
    {
        /* terminate the current thread with status NULL */
        thr_exit(NULL);
    }
    ...
    umem_nofail_callback(my_failure_handler);
    ...
    int my_arg;

    thread_t tid;
    void *status;

    (void) thr_create(NULL, NULL, start_func, &my_arg, 0,
        NULL);
    ...
    while (thr_join(0, &tid, &status) != 0)
        ;

    if (status == NULL) {
        (void) fprintf(stderr, "thread %d ran out of memory\n",
            tid);
    }
    ...

```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	See below.
MT-Level	MT-Safe

umem_alloc(3MALLOC)

The `malloc()`, `calloc()`, `free()`, `realloc()`, and `valloc()` functions are Standard. The `memalign()` function is Stable. The `umem_alloc()`, `umem_zalloc()`, `umem_free()`, and `umem_nofail_callback()` functions are Evolving.

SEE ALSO `exit(2)`, `mmap(2)`, `sbrk(2)`, `bsdmalloc(3MALLOC)`, `libumem(3LIB)`, `longjmp(3C)`, `malloc(3C)`, `malloc(3MALLOC)`, `mapmalloc(3MALLOC)`, `pthread_exit(3C)`, `thr_exit(3C)`, `umem_cache_create(3MALLOC)`, `umem_debug(3MALLOC)`, `watchmalloc(3MALLOC)`, `attributes(5)`, `standards(5)`

Solaris Modular Debugger Guide

WARNINGS Any of the following can cause undefined results:

- Passing a pointer returned from `umem_alloc()` or `umem_zalloc()` to `free()` or `realloc()`.
- Passing a pointer returned from `malloc()`, `calloc()`, `valloc()`, `memalign()`, or `realloc()` to `umem_free()`.
- Writing past the end of a buffer allocated using `umem_alloc()` or `umem_zalloc()`
- Performing `UMEM_NOFAIL` allocations from an `atexit(3C)` handler.

If the `UMEM_NOFAIL` callback performs `UMEM_NOFAIL` allocations, infinite recursion can occur.

NOTES The following list compares the features of the `malloc(3C)`, `bsdmalloc(3MALLOC)`, `malloc(3MALLOC)`, `mtmalloc(3MALLOC)`, and the `libumem` functions.

- The `malloc(3C)`, `bsdmalloc(3MALLOC)`, and `malloc(3MALLOC)` functions have no support for concurrency. The `libumem` and `mtmalloc(3MALLOC)` functions support concurrent allocations.
- The `bsdmalloc(3MALLOC)` functions afford better performance but are space-inefficient.
- The `malloc(3MALLOC)` functions are space-efficient but have slower performance.
- The standard, fully SCD-compliant `malloc(3C)` functions are a trade-off between performance and space-efficiency.
- The `mtmalloc(3MALLOC)` functions provide fast, concurrent `malloc()` implementations that are not space-efficient.
- The `libumem` functions provide a fast, concurrent allocation implementation that in most cases is more space-efficient than `mtmalloc(3MALLOC)`.

NAME	umem_cache_create, umem_cache_destroy, umem_cache_alloc, umem_cache_free – allocation cache manipulation																		
SYNOPSIS	<pre>cc [flag ...] file... -lumem [library ...] #include <umem.h> umem_cache_t *umem_cache_create(char *debug_name, size_t bufsize, size_t align, umem_constructor_t *constructor, umem_destructor_t *destructor, umem_reclaim_t *reclaim, void *callback_data, vmem_t *source, int cflags); void umem_cache_destroy(umem_cache_t *cache); void *umem_cache_alloc(umem_cache_t *cache, int flags); void umem_cache_free(umem_cache_t *cache, void *buffer);</pre>																		
DESCRIPTION	<p>These functions create, destroy, and use an "object cache". An object cache is a collection of buffers of a single size, with optional content caching enabled by the use of callbacks (see Cache Callbacks). Object caches are MT-Safe. Multiple allocations and freeing of memory from different threads can proceed simultaneously. Object caches are faster and use less space per buffer than <code>malloc(3MALLOC)</code> and <code>umem_alloc(3MALLOC)</code>. For more information about object caching, see "The Slab Allocator: An Object-Caching Kernel Memory Allocator" and "Magazines and vmem: Extending the Slab Allocator to Many CPUs and Arbitrary Resources".</p> <p>The <code>umem_cache_create()</code> function creates object caches. Once a cache has been created, objects can be requested from and returned to the cache using <code>umem_cache_alloc()</code> and <code>umem_cache_free()</code>, respectively. A cache with no outstanding buffers can be destroyed with <code>umem_cache_destroy()</code>.</p>																		
Creating and Destroying Caches	<p>The <code>umem_cache_create()</code> function creates a cache of objects and takes as arguments the following:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><i>debug_name</i></td> <td>A human-readable name for debugging purposes.</td> </tr> <tr> <td><i>bufsize</i></td> <td>The size, in bytes, of the buffers in this cache.</td> </tr> <tr> <td><i>align</i></td> <td>The minimum alignment required for buffers in this cache. This parameter must be a power of 2. If 0, it is replaced with the minimum required alignment for the current architecture.</td> </tr> <tr> <td><i>constructor</i></td> <td>The callback to construct an object.</td> </tr> <tr> <td><i>destructor</i></td> <td>The callback to destroy an object.</td> </tr> <tr> <td><i>reclaim</i></td> <td>The callback to reclaim objects.</td> </tr> <tr> <td><i>callback_data</i></td> <td>An opaque pointer passed to the callbacks.</td> </tr> <tr> <td><i>source</i></td> <td>This parameter must be NULL.</td> </tr> <tr> <td><i>cflags</i></td> <td>This parameter must be either 0 or <code>UMC_NODEBUG</code>. If <code>UMC_NODEBUG</code>, all debugging features are disabled for this cache. See <code>umem_debug(3MALLOC)</code>.</td> </tr> </table>	<i>debug_name</i>	A human-readable name for debugging purposes.	<i>bufsize</i>	The size, in bytes, of the buffers in this cache.	<i>align</i>	The minimum alignment required for buffers in this cache. This parameter must be a power of 2. If 0, it is replaced with the minimum required alignment for the current architecture.	<i>constructor</i>	The callback to construct an object.	<i>destructor</i>	The callback to destroy an object.	<i>reclaim</i>	The callback to reclaim objects.	<i>callback_data</i>	An opaque pointer passed to the callbacks.	<i>source</i>	This parameter must be NULL.	<i>cflags</i>	This parameter must be either 0 or <code>UMC_NODEBUG</code> . If <code>UMC_NODEBUG</code> , all debugging features are disabled for this cache. See <code>umem_debug(3MALLOC)</code> .
<i>debug_name</i>	A human-readable name for debugging purposes.																		
<i>bufsize</i>	The size, in bytes, of the buffers in this cache.																		
<i>align</i>	The minimum alignment required for buffers in this cache. This parameter must be a power of 2. If 0, it is replaced with the minimum required alignment for the current architecture.																		
<i>constructor</i>	The callback to construct an object.																		
<i>destructor</i>	The callback to destroy an object.																		
<i>reclaim</i>	The callback to reclaim objects.																		
<i>callback_data</i>	An opaque pointer passed to the callbacks.																		
<i>source</i>	This parameter must be NULL.																		
<i>cflags</i>	This parameter must be either 0 or <code>UMC_NODEBUG</code> . If <code>UMC_NODEBUG</code> , all debugging features are disabled for this cache. See <code>umem_debug(3MALLOC)</code> .																		

umem_cache_create(3MALLOC)

Each cache can have up to three associated callbacks:

```
int constructor(void *buffer, void *callback_data, int flags);
void destructor(void *buffer, void *callback_data);
void reclaim(void *callback_data);
```

The *callback_data* argument is always equal to the value passed to `umem_cache_create()`, thereby allowing a client to use the same callback functions for multiple caches, but with customized behavior.

The reclaim callback is called when the `umem` function is requesting more memory from the operating system. This callback can be used by clients who retain objects longer than they are strictly needed (for example, caching non-active state). A typical reclaim callback might return to the cache ten per cent of the unneeded buffers.

The constructor and destructor callbacks enable the management of buffers with the constructed state. The constructor takes as arguments a buffer with undefined contents, some callback data, and the flags to use for any allocations. This callback should transform the buffer into the constructed state.

The destructor callback takes as an argument a constructed object and prepares it for return to the general pool of memory. The destructor should undo any state that the constructor created. For debugging, the destructor can also check that the buffer is in the constructed state, to catch incorrectly freed buffers. See `umem_debug(3MALLOC)` for further information on debugging support.

The `umem_cache_destroy()` function destroys an object cache. If the cache has any outstanding allocations, the behavior is undefined.

Allocating Objects

The `umem_cache_alloc()` function takes as arguments:

<i>cache</i>	a cache pointer
<i>flags</i>	flags that determine the behavior if <code>umem_cache_alloc()</code> is unable to fulfill the allocation request

If successful, `umem_cache_alloc()` returns a pointer to the beginning of an object of *bufsize* length.

There are three cases to consider:

- A new buffer needed to be allocated. If the cache was created with a constructor, it is applied to the buffer and the resulting object is returned.
- The object cache was able to use a previously freed buffer. If the cache was created with a constructor, the object is returned unchanged from when it was freed.
- The allocation of a new buffer failed. The *flags* argument determines the behavior:

MEM_DEFAULT	The <code>umem_cache_alloc()</code> function returns NULL if the allocation fails.
-------------	--

umem_cache_create(3MALLOC)

UMEM_NOFAIL The `umem_cache_alloc()` function cannot return NULL. A callback is used to determine what action occurs. See `umem_alloc(3MALLOC)` for more information.

Freeing Objects

The `umem_cache_free()` function takes as arguments:

cache a cache pointer

buf a pointer previously returned from `umem_cache_alloc()`. This argument must not be NULL.

If the cache was created with a constructor callback, the object must be returned to the constructed state before it is freed.

Undefined behavior results if an object is freed multiple times, if an object is modified after it is freed, or if an object is freed to a cache other than the one from which it was allocated.

Caches with Constructors

When a constructor callback is in use, there is essentially a contract between the cache and its clients. The cache guarantees that all objects returned from `umem_cache_alloc()` will be in the constructed state, and the client guarantees that it will return the object to the constructed state before handing it to `umem_cache_free()`.

RETURN VALUES

Upon failure, the `umem_cache_create()` function returns a null pointer.

ERRORS

The `umem_cache_create()` function will fail if:

EAGAIN There is not enough memory available to allocate the cache data structure.

EINVAL The *debug_name* argument is NULL, the *align* argument is not a power of two or is larger than the system *pagesize*, or the *bufsize* argument is 0.

ENOMEM The `libumem` library could not be initialized, or the *bufsize* argument is too large and its use would cause integer overflow to occur.

EXAMPLES

EXAMPLE 1 Use a fixed-size structure with no constructor callback.

```
#include <umem.h>

typedef struct my_obj {
    long my_data1;
} my_obj_t;

/*
 * my_objs can be freed at any time. The contents of
 * my_data1 is undefined at allocation time.
 */

umem_cache_t *my_obj_cache;
```

umem_cache_create(3MALLOC)

EXAMPLE 1 Use a fixed-size structure with no constructor callback. (Continued)

```
...
my_obj_cache = umem_cache_create("my_obj", sizeof (my_obj_t),
    0, NULL, NULL, NULL, NULL, NULL, 0);
...
my_obj_t *cur = umem_cache_alloc(my_obj_cache, UMEM_DEFAULT);
...
/* use cur */
...
umem_cache_free(my_obj_cache, cur);
...
```

EXAMPLE 2 Use an object with a mutex.

```
#define _REENTRANT
#include <synch.h>
#include <umem.h>

typedef struct my_obj {
    mutex_t my_mutex;
    long my_data;
} my_obj_t;

/*
 * my_objs can only be freed when my_mutex is unlocked.
 */
int
my_obj_constructor(void *buf, void *ignored, int flags)
{
    my_obj_t *myobj = buf;

    (void) mutex_init(&my_obj->my_mutex, USYNC_THREAD, NULL);

    return (0);
}

void
my_obj_destructor(void *buf, void *ignored)
{
    my_obj_t *myobj = buf;

    (void) mutex_destroy(&my_obj->my_mutex);
}

umem_cache_t *my_obj_cache;

...
my_obj_cache = umem_cache_create("my_obj", sizeof (my_obj_t),
    0, my_obj_constructor, my_obj_destructor, NULL, NULL,
    NULL, 0);
...
my_obj_t *cur = umem_cache_alloc(my_obj_cache, UMEM_DEFAULT);
cur->my_data = 0;          /* cannot assume anything about my_data */
...
```

EXAMPLE 2 Use an object with a mutex. (Continued)

```
umem_cache_free(my_obj_cache, cur);
...
```

EXAMPLE 3 Use a more complex object with a mutex.

```
#define _REENTRANT
#include <assert.h>
#include <synch.h>
#include <umem.h>

typedef struct my_obj {
    mutex_t my_mutex;
    cond_t my_cv;
    struct bar *my_barlist;
    unsigned my_refcount;
} my_obj_t;

/*
 * my_objs can only be freed when my_barlist == NULL,
 * my_refcount == 0, there are no waiters on my_cv, and
 * my_mutex is unlocked.
 */

int
my_obj_constructor(void *buf, void *ignored, int flags)
{
    my_obj_t *myobj = buf;

    (void) mutex_init(&my_obj->my_mutex, USYNC_THREAD, NULL);
    (void) cond_init(&my_obj->my_cv, USYNC_THREAD, NULL);
    myobj->my_barlist = NULL;
    myobj->my_refcount = 0;

    return (0);
}

void
my_obj_destructor(void *buf, void *ignored)
{
    my_obj_t *myobj = buf;

    assert(myobj->my_refcount == 0);
    assert(myobj->my_barlist == NULL);
    (void) cond_destroy(&my_obj->my_cv);
    (void) mutex_destroy(&my_obj->my_mutex);
}

umem_cache_t *my_obj_cache;

...
my_obj_cache = umem_cache_create("my_obj", sizeof (my_obj_t),
    0, my_obj_constructor, my_obj_destructor, NULL, NULL,
    NULL, 0);
...
```

umem_cache_create(3MALLOC)

EXAMPLE 3 Use a more complex object with a mutex. (Continued)

```
my_obj_t *cur = umem_cache_alloc(my_obj_cache, UMEM_DEFAULT);
...
/* use cur */
...
umem_cache_free(my_obj_cache, cur);
...
```

EXAMPLE 4 Use objects with a subordinate buffer while reusing callbacks.

```
#include assert.h>
#include umem.h>

typedef struct my_obj {
    char *my_buffer;
    size_t my_size;
} my_obj_t;

/*
 * my_size and the my_buffer pointer should never be changed
 */

int
my_obj_constructor(void *buf, void *arg, int flags)
{
    size_t sz = (size_t)arg;

    my_obj_t *myobj = buf;

    if ((myobj->my_buffer = umem_alloc(sz, flags)) == NULL)
        return (1);

    my_size = sz;

    return (0);
}

void
my_obj_destructor(void *buf, void *arg)
{
    size_t sz = (size_t)arg;

    my_obj_t *myobj = buf;

    assert(sz == buf->my_size);
    umem_free(myobj->my_buffer, sz);
}

...
umem_cache_t *my_obj_4k_cache;
umem_cache_t *my_obj_8k_cache;
...
my_obj_cache_4k = umem_cache_create("my_obj_4k", sizeof (my_obj_t),
    0, my_obj_constructor, my_obj_destructor, NULL, (void *)4096,
    NULL, 0);
```


EXAMPLE 4 Use objects with a subordinate buffer while reusing callbacks. (Continued)

```
my_obj_cache_8k = umem_cache_create("my_obj_8k", sizeof (my_obj_t),
    0, my_obj_constructor, my_obj_destructor, NULL, (void *)8192,
    NULL, 0);
...
my_obj_t *my_obj_4k = umem_cache_alloc(my_obj_4k_cache,
    UMEM_DEFAULT);
my_obj_t *my_obj_8k = umem_cache_alloc(my_obj_8k_cache,
    UMEM_DEFAULT);
/* no assumptions should be made about the contents of the buffers */
...
/* make sure to return them to the correct cache */
umem_cache_free(my_obj_4k_cache, my_obj_4k);
umem_cache_free(my_obj_8k_cache, my_obj_8k);
...
```

See the **EXAMPLES** section of [umem_alloc\(3MALLOC\)](#) for examples involving the `UMEM_NOFAIL` flag.

ATTRIBUTES See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Evolving
MT-Level	MT-Safe

SEE ALSO [setcontext\(2\)](#), [atexit\(3C\)](#), [libumem\(3LIB\)](#), [longjmp\(3C\)](#), [swapcontext\(3C\)](#), [thr_exit\(3C\)](#), [umem_alloc\(3MALLOC\)](#), [umem_debug\(3MALLOC\)](#), [attributes\(5\)](#)

Bonwick, Jeff, "The Slab Allocator: An Object-Caching Kernel Memory Allocator", Proceedings of the Summer 1994 Usenix Conference.

Bonwick, Jeff and Jonathan Adams, "Magazines and vmem: Extending the Slab Allocator to Many CPUs and Arbitrary Resources", Proceedings of the Summer 2001 Usenix Conference.

WARNINGS Any of the following can cause undefined results:

- Destroying a cache that has outstanding allocated buffers.
- Using a cache after it has been destroyed.
- Calling `umem_cache_free()` on the same buffer multiple times.
- Passing a `NULL` pointer to `umem_cache_free()`.
- Writing past the end of a buffer.
- Reading from or writing to a buffer after it has been freed.
- Performing `UMEM_NOFAIL` allocations from an [atexit\(3C\)](#) handler.

umem_cache_create(3MALLOC)

Per-cache callbacks can be called from a variety of contexts. The use of functions that modify the active context, such as `setcontext(2)`, `swapcontext(3C)`, and `thr_exit(3C)`, or functions that are unsafe for use in multithreaded applications, such as `longjmp(3C)` and `siglongjmp(3C)`, result in undefined behavior.

A constructor callback that performs allocations must pass its *flags* argument unchanged to `umem_alloc(3MALLOC)` and `umem_cache_alloc()`. Any allocations made with a different *flags* argument results in undefined behavior. The constructor must correctly handle the failure of any allocations it makes.

NOTES Object caches make the following guarantees about objects:

- If the cache has a constructor callback, it is applied to every object before it is returned from `umem_cache_alloc()` for the first time.
- If the cache has a constructor callback, an object passed to `umem_cache_free()` and later returned from `umem_cache_alloc()` is not modified between the two events.
- If the cache has a destructor, it is applied to all objects before their underlying storage is returned.

No other guarantees are made. In particular, even if there are buffers recently freed to the cache, `umem_cache_alloc()` can fail.

NAME	umem_debug – debugging features of the umem library
SYNOPSIS	<pre>cc [<i>flag...</i>] <i>file...</i> -lumem [<i>library...</i>] #include <umem.h></pre>
DESCRIPTION	<p>The <code>libumem</code> library provides debugging features that detect memory leaks, buffer overruns, multiple frees, use of uninitialized data, use of freed data, and many other common programming errors. The activation of the run-time debugging features is controlled by environment variables.</p> <p>When the library detects an error, it writes a description of the error to an internal buffer that is readable with the <code>::umem_status mdb(1) dcmd</code> and then calls <code>abort(3C)</code>.</p>
ENVIRONMENT VARIABLES	<p>UMEM_DEBUG This variable contains a list of comma-separated options. Unrecognized options are ignored. Possible options include:</p> <p>audit[=<i>frames</i>] This option enables the recording of auditing information, including thread ID, high-resolution time stamp, and stack trace for the last action (allocation or free) on every allocation. If transaction logging (see <code>UMEM_LOGGING</code>) is enabled, this auditing information is also logged.</p> <p>The <i>frames</i> parameter sets the number of stack frames recorded in the auditing structure. The upper bound for frames is implementation-defined. If a larger value is requested, the upper bound is used instead.</p> <p>If <i>frames</i> is not specified or is not an integer, the default value of 15 is used.</p> <p>This option also enables the <code>guards</code> option.</p> <p>contents[=<i>count</i>] If auditing and contents logging (see <code>UMEM_LOGGING</code>) are enabled, the first <i>count</i> bytes of each buffer are logged when they are freed. If a buffer is shorter than <i>count</i> bytes, it is logged in its entirety.</p> <p>If <i>count</i> is not specified or is not an integer, the default value of 256 is used.</p> <p>default This option is equivalent to <code>audit,contents,guards</code>.</p> <p>guards This option enables filling allocated and freed buffers with special patterns to help detect the use of uninitialized data and previously freed buffers. It also enables an 8-byte redzone after each buffer that contains <code>0xfeedfacefeedfaceULL</code>.</p>

umem_debug(3MALLOC)

When an object is freed, it is filled with `0xdeadbeef`. When an object is allocated, the `0xdeadbeef` pattern is verified and replaced with `0xbaddcafe`. The redzone is checked every time a buffer is allocated or freed.

For caches with either constructors or destructors, or both, `umem_cache_alloc(3MALLOC)` and `umem_cache_free(3MALLOC)` apply the cache's constructor and destructor, respectively, instead of caching constructed objects. The presence of `assert(3C)`s in the destructor verifying that the buffer is in the constructed state can be used to detect any objects returned in an improper state. See `umem_cache_create(3MALLOC)` for details.

verbose

The library writes error descriptions to standard error before aborting. These messages are not localized.

UMEM_LOGGING

To be enabled, this variable should be set to a comma-separated list of in-memory logs. The logs available are:

transaction[=*size*]

If the `audit` debugging option is set (see `UMEM_DEBUG`), the audit structures from previous transactions are entered into this log.

contents[=*size*]

If the `audit` debugging option is set, the contents of objects are recorded in this log as they are freed.

If the "contents" debugging option was not set, 256 bytes of each freed buffer are saved.

fail[=*size*]

Records are entered into this log for every failed allocation.

For any of these options, if *size* is not specified, the default value of 64k is used. The *size* parameter must be an integer that can be qualified with K, M, G, or T to specify kilobytes, megabytes, gigabytes, or terabytes, respectively.

Logs that are not listed or that have either a size of 0 or an invalid size are disabled.

The log is disabled if during initialization the requested amount of storage cannot be allocated.

umem_debug(3MALLOC)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Unstable
MT-Level	MT-Safe

SEE ALSO `mdb(1)`, `abort(3C)`, `signal(3C)`, `umem_cache_create(3MALLOC)`, `attributes(5)`

Solaris Modular Debugger Guide

WARNINGS When `libumem` aborts the process using `abort(3C)`, any existing signal handler for `SIGABRT` is called. If the signal handler performs allocations, undefined behavior can result.

NOTES Some of the debugging features work only for allocations smaller than 16 kilobytes in size. Allocations larger than 16 kilobytes could have reduced support.

Activating any of the library's debugging features could significantly increase the library's memory footprint and decrease its performance.

ungetc(3C)

NAME	ungetc – push byte back into input stream						
SYNOPSIS	<pre>#include <stdio.h> int ungetc(int c, FILE *stream);</pre>						
DESCRIPTION	<p>The <code>ungetc()</code> function pushes the byte specified by <code>c</code> (converted to an unsigned char) back onto the input stream pointed to by <code>stream</code>. The pushed-back bytes will be returned by subsequent reads on that stream in the reverse order of their pushing. A successful intervening call (with the stream pointed to by <code>stream</code>) to a file-positioning function (<code>fseek(3C)</code>, <code>fsetpos(3C)</code> or <code>rewind(3C)</code>) discards any pushed-back bytes for the stream. The external storage corresponding to the stream is unchanged.</p> <p>Four bytes of push-back are guaranteed. If <code>ungetc()</code> is called too many times on the same stream without an intervening read or file-positioning operation on that stream, the operation may fail.</p> <p>If the value of <code>c</code> equals that of the macro <code>EOF</code>, the operation fails and the input stream is unchanged.</p> <p>A successful call to <code>ungetc()</code> clears the end-of-file indicator for the stream. The value of the file-position indicator for the stream after reading or discarding all pushed-back bytes will be the same as it was before the bytes were pushed back. The file-position indicator is decremented by each successful call to <code>ungetc()</code>; if its value was 0 before a call, its value is indeterminate after the call.</p>						
RETURN VALUES	Upon successful completion, <code>ungetc()</code> returns the byte pushed back after conversion. Otherwise it returns <code>EOF</code> .						
ERRORS	No errors are defined.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>read(2)</code> , <code>intro(3)</code> , <code>__fsetlocking(3C)</code> , <code>fseek(3C)</code> , <code>fsetpos(3C)</code> , <code>getc(3C)</code> , <code>setbuf(3C)</code> , <code>stdio(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	ungetwc – push wide-character code back into input stream						
SYNOPSIS	<pre>#include <stdio.h> #include <wchar.h> wint_t ungetwc(wint_t wc, FILE *stream);</pre>						
DESCRIPTION	<p>The <code>ungetwc()</code> function pushes the character corresponding to the wide character code specified by <code>wc</code> back onto the input stream pointed to by <code>stream</code>. The pushed-back characters will be returned by subsequent reads on that stream in the reverse order of their pushing. A successful intervening call (with the stream pointed to by <code>stream</code>) to a file-positioning function (<code>fseek(3C)</code>, <code>fsetpos(3C)</code> or <code>rewind(3C)</code>) discards any pushed-back characters for the stream. The external storage corresponding to the stream is unchanged.</p> <p>One character of push-back is guaranteed. If <code>ungetwc()</code> is called too many times on the same stream without an intervening read or file-positioning operation on that stream, the operation may fail.</p> <p>If the value of <code>wc</code> equals that of the macro <code>WEOF</code>, the operation fails and the input stream is unchanged.</p> <p>A successful call to <code>ungetwc()</code> clears the end-of-file indicator for the stream. The value of the file-position indicator for the stream after reading or discarding all pushed-back characters will be the same as it was before the characters were pushed back. The file-position indicator is decremented (by one or more) by each successful call to <code>ungetwc()</code>; if its value was 0 before a call, its value is indeterminate after the call.</p>						
RETURN VALUES	Upon successful completion, <code>ungetwc()</code> returns the wide-character code corresponding to the pushed-back character. Otherwise it returns <code>WEOF</code> .						
ERRORS	<p>The <code>ungetwc()</code> function may fail if:</p> <p>EILSEQ An invalid character sequence is detected, or a wide-character code does not correspond to a valid character.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>read(2)</code> , <code>fseek(3C)</code> , <code>fsetpos(3C)</code> , <code>rewind(3C)</code> , <code>setbuf(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

unlockpt(3C)

NAME unlockpt – unlock a pseudo-terminal master/slave pair

SYNOPSIS

```
#include <stdlib.h>
int unlockpt (int fildev);
```

DESCRIPTION The `unlockpt()` function unlocks the slave pseudo-terminal device associated with the master to which *fildev* refers.

Portable applications must call `unlockpt()` before opening the slave side of a pseudo-terminal device.

RETURN VALUES Upon successful completion, `unlockpt()` returns 0. Otherwise, it returns -1 and sets `errno` to indicate the error.

ERRORS The `unlockpt()` function may fail if:

EBADF	The <i>fildev</i> argument is not a file descriptor open for writing.
EINVAL	The <i>fildev</i> argument is not associated with a master pseudo-terminal device.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	Safe

SEE ALSO `open(2)`, `grantpt(3C)`, `ptsname(3C)`, `attributes(5)`, `standards(5)`
STREAMS Programming Guide

NAME	unsetenv – remove an environment variable						
SYNOPSIS	<pre>#include <stdlib.h> int unsetenv(const char *name);</pre>						
DESCRIPTION	<p>The <code>unsetenv()</code> function removes an environment variable from the environment of the calling process. The <i>name</i> argument points to a string that is the name of the variable to be removed. The named argument cannot contain an '=' character. If the named variable does not exist in the current environment, the environment is unchanged and the function is considered to have completed successfully.</p> <p>If the application modifies <i>environ</i> or the pointers to which it points, the behavior of <code>unsetenv()</code> is undefined. The <code>unsetenv()</code> function updates the list of pointers to which <i>environ</i> points.</p>						
RETURN VALUES	Upon successful completion, 0 is returned. Otherwise, -1 is returned, <code>errno</code> set to indicate the error, and the environment is left unchanged.						
ERRORS	<p>The <code>unsetenv()</code> function will fail if:</p> <p>EINVAL The <i>name</i> argument is a null pointer, points to an empty string, or points to a string containing an '=' character.</p>						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>getenv(3C)</code> , <code>setenv(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

usleep(3C)

NAME	usleep – suspend execution for interval in microseconds						
SYNOPSIS	<pre>#include <unistd.h> int usleep(useconds_t <i>useconds</i>);</pre>						
DESCRIPTION	<p>The <code>usleep()</code> function suspends the caller from execution for the number of microseconds specified by the <i>useconds</i> argument. (A microsecond is .000001 seconds.) Because of other activity, or because of the time spent in processing the call, the actual suspension time may be longer than the amount of time specified.</p> <p>If the value of <i>useconds</i> is 0, then the call has no effect.</p> <p>In a single-threaded program (one not linked with <code>-lthread</code> or <code>-lpthread</code>), the <code>usleep()</code> function uses the process's realtime interval timer to indicate to the system when the process should be woken up.</p> <p>There is one real-time interval timer for each process. The <code>usleep()</code> function will not interfere with a previous setting of this timer. If the process has set this timer prior to calling <code>usleep()</code>, and if the time specified by <i>useconds</i> equals or exceeds the interval timer's prior setting, the caller will be woken up shortly before the timer was set to expire.</p> <p>Interactions between <code>usleep()</code> and either <code>alarm(2)</code> or <code>sleep(3C)</code> are unspecified.</p> <p>In a multithreaded program (one linked with <code>-lthread</code> or <code>-lpthread</code>), <code>usleep()</code> is implemented by a call to <code>nanosleep(3RT)</code> and does not modify the state of the alarm signal or the realtime interval timer. There is no interaction between this version of <code>usleep()</code> and either <code>alarm(2)</code> or <code>sleep(3C)</code>.</p>						
RETURN VALUES	On completion, <code>usleep()</code> returns 0. There are no error retruns.						
ERRORS	No errors are returned.						
USAGE	The <code>usleep()</code> function is included for its historical usage. The <code>nanosleep(3RT)</code> function is preferred over this function.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	Safe						
SEE ALSO	<code>alarm(2)</code> , <code>setitimer(2)</code> , <code>sigaction(2)</code> , <code>sigprocmask(2)</code> , <code>nanosleep(3RT)</code> , <code>sleep(3C)</code> , <code>ualarm(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						
NOTES	In a multithreaded program, only the invoking thread is suspended from execution.						

NAME	vfwprintf, vswprintf, vwprintf – wide-character formatted output of a stdarg argument list						
SYNOPSIS	<pre>#include <stdarg.h> #include <stdio.h> #include <wchar.h> int vfwprintf(FILE *restrict <i>stream</i>, const wchar_t *restrict <i>format</i>, va_list <i>arg</i>); int vswprintf(wchar_t *restrict <i>s</i>, size_t <i>n</i>, const wchar_t *restrict <i>format</i>, va_list <i>arg</i>); int vwprintf(const wchar_t *restrict <i>format</i>, va_list <i>arg</i>);</pre>						
DESCRIPTION	<p>The vwprintf(), vfwprintf(), and vswprintf() functions are the same as wprintf(), fwprintf(), and swprintf() respectively, except that instead of being called with a variable number of arguments, they are called with an argument list as defined by <stdarg.h>.</p> <p>These functions do not invoke the va_end() macro. However, as these functions do invoke the va_arg() macro, the value of <i>ap</i> after the return is indeterminate.</p>						
RETURN VALUES	Refer to fwprintf(3C) .						
ERRORS	Refer to fwprintf(3C) .						
USAGE	Applications using these functions should call va_end(<i>ap</i>) afterwards to clean up.						
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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe with exceptions						
SEE ALSO	fwprintf(3C) , setlocale(3C) , attributes(5) , standards(5)						
NOTES	The vwprintf(), vfwprintf(), and vswprintf() functions can be used safely in multithreaded applications, as long as setlocale(3C) is not being called to change the locale.						

vlfmt(3C)

NAME	vlfmt – display error message in standard format and pass to logging and monitoring services
SYNOPSIS	<pre>#include <pfmt.h> #include <stdarg.h> int vlfmt(FILE *stream, long flag, const char *format, va_list ap);</pre>
DESCRIPTION	<p>The vlfmt () function is identical to lfmt(3C), except that it is called with an argument list as defined by <stdarg.h>.</p> <p>The <stdarg.h> header defines the type va_list and a set of macros for advancing through a list of arguments whose number and types may vary. The ap argument is of type va_list. This argument is used with the <stdarg.h> macros va_start (), va_arg (), and va_end (). See stdarg(3EXT). The example in the EXAMPLES section below demonstrates their use with vlfmt ().</p>
RETURN VALUES	Upon successful completion, vlfmt () returns the number of bytes transmitted. Otherwise, -1 is returned if there was a write error to stream, or -2 is returned if unable to log and/or display at console.
EXAMPLES	<p>EXAMPLE 1 Use of vlfmt () to write an errlog () routine.</p> <p>The following example demonstrates how vlfmt () could be used to write an errlog () routine. The va_alist () macro is used as the parameter list in a function definition. The va_start (ap, ...) call, where ap is of type va_list, must be invoked before any attempt to traverse and access unnamed arguments. Calls to va_arg (ap, atype) traverse the argument list. Each execution of va_arg () expands to an expression with the value and type of the next argument in the list ap, which is the same object initialized by va_start (). The atype argument is the type that the returned argument is expected to be. The va_end (ap) macro must be invoked when all desired arguments have been accessed. The argument list in ap can be traversed again if va_start () is called again after va_end (). In the example below, va_arg () is executed first to retrieve the format string passed to errlog (). The remaining errlog () arguments (arg1, arg2, ...) are passed to vlfmt () in the argument ap.</p> <pre>#include <pfmt.h> #include <stdarg.h> /* * errlog should be called like * errlog(log_info, format, arg1, ...); */ void errlog(long log_info, ...) { va_list ap; char *format; va_start(ap,); format = va_arg(ap, char *); (void) vlfmt(stderr, log_info MM_ERROR, format, ap); va_end(ap); (void) abort(); }</pre>

EXAMPLE 1 Use of `vlfmt()` to write an `errlog()` routine. *(Continued)*

USAGE Since `vlfmt()` uses `gettxt(3C)`, it is recommended that `vlfmt()` not be used.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `gettxt(3C)`, `lfmt(3C)`, `attributes(5)`, `stdarg(3EXT)`

vpfmt(3C)

NAME	vpfmt – display error message in standard format and pass to logging and monitoring services
SYNOPSIS	<pre>#include <pfmt.h> #include <stdarg.h> int vpfmt(FILE *stream, long flag, const char *format, va_list ap);</pre>
DESCRIPTION	<p>The <code>vpfmt()</code> function is identical to <code>pfmt(3C)</code>, except that it is called with an argument list as defined by <code><stdarg.h></code>.</p> <p>The <code><stdarg.h></code> header defines the type <code>va_list</code> and a set of macros for advancing through a list of arguments whose number and types may vary. The <code>ap</code> argument is of type <code>va_list</code>. This argument is used with the <code><stdarg.h></code> macros <code>va_start()</code>, <code>va_arg()</code>, and <code>va_end()</code>. See <code>stdarg(3EXT)</code>. The example in the EXAMPLES section below demonstrates their use with <code>vpfmt()</code>.</p>
RETURN VALUES	Upon successful completion, <code>vpfmt()</code> returns the number of bytes transmitted. Otherwise, <code>-1</code> is returned if there was a write error to <code>stream</code> .
EXAMPLES	<p>EXAMPLE 1 Use of <code>vpfmt()</code> to write an error routine.</p> <p>The following example demonstrates how <code>vpfmt()</code> could be used to write an <code>error()</code> routine. The <code>va_alist()</code> macro is used as the parameter list in a function definition. The <code>va_start(ap, ...)</code> call, where <code>ap</code> is of type <code>va_list</code>, must be invoked before any attempt to traverse and access unnamed arguments. Calls to <code>va_arg(ap, atype)</code> traverse the argument list. Each execution of <code>va_arg()</code> expands to an expression with the value and type of the next argument in the list <code>ap</code>, which is the same object initialized by <code>va_start()</code>. The <code>atype</code> argument is the type that the returned argument is expected to be. The <code>va_end(ap)</code> macro must be invoked when all desired arguments have been accessed. The argument list in <code>ap</code> can be traversed again if <code>va_start()</code> is called again after <code>va_end()</code>. In the example below, <code>va_arg()</code> is executed first to retrieve the format string passed to <code>error()</code>. The remaining <code>error()</code> arguments (<code>arg1, arg2, ...</code>) are passed to <code>vpfmt()</code> in the argument <code>ap</code>.</p> <pre>#include <pfmt.h> #include <stdarg.h> /* * error should be called like * error(format, arg1, ...); */ void error(...) { va_list ap; char *format; va_start(ap,); format = va_arg(ap, char *); (void) vpfmt(stderr, MM_ERROR, format, ap); va_end(ap); (void) abort(); }</pre>
USAGE	Since <code>vpfmt()</code> uses <code>gettext(3C)</code> , it is recommended that <code>vpfmt()</code> not be used.

vpfmt(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `gettxt(3C)`, `pfmt(3C)`, `attributes(5)`, `stdarg(3EXT)`

vprintf(3C)

NAME	vprintf, vfprintf, vsprintf, vsnprintf – print formatted output of a variable argument list
SYNOPSIS	<pre>#include <stdio.h> #include <stdarg.h> int vprintf(const char *format, va_list ap); int vfprintf(FILE *stream, const char *format, va_list ap); int vsprintf(char *s, const char *format, va_list ap); int vsnprintf(char *s, size_t n, const char *format, va_list ap);</pre>
DESCRIPTION	<p>The vprintf(), vfprintf(), vsprintf() and vsnprintf() functions are the same as printf(), fprintf(), sprintf(), and snprintf(), respectively, except that instead of being called with a variable number of arguments, they are called with an argument list as defined in the <stdarg.h> header. See printf(3C).</p> <p>The <stdarg.h> header defines the type va_list and a set of macros for advancing through a list of arguments whose number and types may vary. The argument ap to the vprint family of functions is of type va_list. This argument is used with the <stdarg.h> header file macros va_start(), va_arg(), and va_end() (see stdarg(3EXT)). The EXAMPLES section below demonstrates the use of va_start() and va_end() with vprintf().</p> <p>The macro va_alist() is used as the parameter list in a function definition, as in the function called error() in the example below. The macro va_start(ap, parmN), where ap is of type va_list and parmN is the rightmost parameter (just before ...), must be called before any attempt to traverse and access unnamed arguments is made. The va_end(ap) macro must be invoked when all desired arguments have been accessed. The argument list in ap can be traversed again if va_start() is called again after va_end(). In the example below, the error() arguments (arg1, arg2, ...) are passed to vfprintf() in the argument ap.</p>
RETURN VALUES	Refer to printf(3C).
ERRORS	The vprintf() and vfprintf() functions will fail if either the stream is unbuffered or the stream's buffer needed to be flushed and: EFBIG The file is a regular file and an attempt was made to write at or beyond the offset maximum.
EXAMPLES	<p>EXAMPLE 1 Using vprintf() to write an error routine.</p> <p>The following demonstrates how vfprintf() could be used to write an error routine:</p> <pre>#include <stdio.h> #include <stdarg.h> ... /* * error should be called like * error(function_name, format, arg1, ...);</pre>

EXAMPLE 1 Using `vprintf()` to write an error routine. (Continued)

```

*/
void error(char *function_name, char *format, ...)
{
    va_list ap;
    va_start(ap, format);
    /* print out name of function causing error */
    (void) fprintf(stderr, "ERR in %s: ", function_name);
    /* print out remainder of message */
    (void) vfprintf(stderr, format, ap);
    va_end(ap);
    (void) abort();
}

```

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	See below.

The `vprintf()` and `vfprintf()` functions are MT-Safe. The `vsprintf()` and `vsnprintf()` functions are Async-Signal-Safe.

SEE ALSO [printf\(3C\)](#), [attributes\(5\)](#), [stdarg\(3EXT\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NOTES The `vsprintf()` return value when $n = 0$ was changed in the Solaris 10 release. The change was based on the SUSv3 specification. The previous behavior was based on the initial SUSv2 specification, where `vsprintf()` when $n = 0$ returns an unspecified value less than 1.

vsyslog(3C)

- NAME** vsyslog – log message with a stdarg argument list
- SYNOPSIS**
- ```
#include <syslog.h>
#include <stdarg.h>

int vsyslog(int priority, const char *message, va_list ap);
```
- DESCRIPTION** The vsyslog() function is identical to [syslog\(3C\)](#), except that it is called with an argument list as defined by <stdarg.h> rather than with a variable number of arguments.
- EXAMPLES**
- EXAMPLE 1** Use vsyslog() to write an error routine.
- The following demonstrates how vsyslog() can be used to write an error routine.
- ```
#include <syslog.h>
#include <stdarg.h>

/*
 * error should be called like:
 * error(pri, function_name, format, arg1, arg2...);
 */

void
error(int pri, char *function_name, char *format, ...)
{
    va_list args;

    va_start(args, format);
    /* log name of function causing error */
    (void) syslog(pri, "ERROR in %s.", function_name);
    /* log remainder of message */
    (void) vsyslog(pri, format, args);
    va_end(args);
    (void) abort();
}

main()
{
    error(LOG_ERR, "main", "process %d is dying", getpid());
}
```
- ATTRIBUTES** See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	Safe

SEE ALSO [syslog\(3C\)](#), [attributes\(5\)](#)

NAME	wait3, wait4 – wait for process to terminate or stop				
SYNOPSIS	<pre>#include <sys/wait.h> #include <sys/time.h> #include <sys/resource.h> pid_t wait3(int *statusp, int options, struct rusage *rusage); pid_t wait4(pid_t pid, int *statusp, int options, struct rusage *rusage);</pre>				
DESCRIPTION	<p>The <code>wait3()</code> function delays its caller until a signal is received or one of its child processes terminates or stops due to tracing. If any child process has died or stopped due to tracing and this has not already been reported, return is immediate, returning the process ID and status of one of those children. If that child process has died, it is discarded. If there are no children, <code>-1</code> is returned immediately. If there are only running or stopped but reported children, the calling process is blocked.</p> <p>If <code>statusp</code> is not a null pointer, then on return from a successful <code>wait3()</code> call, the status of the child process is stored in the integer pointed to by <code>statusp</code>. <code>*statusp</code> indicates the cause of termination and other information about the terminated process in the following manner:</p> <ul style="list-style-type: none"> ■ If the low-order 8 bits of <code>*statusp</code> are equal to <code>0177</code>, the child process has stopped; the 8 bits higher up from the low-order 8 bits of <code>*statusp</code> contain the number of the signal that caused the process to stop. See <code>signal.h(3HEAD)</code>. ■ If the low-order 8 bits of <code>*statusp</code> are non-zero and are not equal to <code>0177</code>, the child process terminated due to a signal; the low-order 7 bits of <code>*statusp</code> contain the number of the signal that terminated the process. In addition, if the low-order seventh bit of <code>*statusp</code> (that is, bit <code>0200</code>) is set, a “core image” of the process was produced; see <code>signal.h(3HEAD)</code>. ■ Otherwise, the child process terminated due to an <code>exit()</code> call; the 8 bits higher up from the low-order 8 bits of <code>*statusp</code> contain the low-order 8 bits of the argument that the child process passed to <code>exit()</code>; see <code>exit(2)</code>. <p>The <code>options</code> argument is constructed from the bitwise inclusive OR of zero or more of the following flags, defined in <code><sys/wait.h></code>:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;"><code>WNOHANG</code></td> <td>Execution of the calling process is not suspended if status is not immediately available for any child process.</td> </tr> <tr> <td><code>WUNTRACED</code></td> <td>The status of any child processes that are stopped, and whose status has not yet been reported since they stopped, are also reported to the requesting process.</td> </tr> </table> <p>If <code>rusage</code> is not a null pointer, a summary of the resources used by the terminated process and all its children is returned. Only the user time used and the system time used are currently available. They are returned in the <code>ru_utime</code> and <code>ru_stime</code>, members of the <code>rusage</code> structure, respectively.</p>	<code>WNOHANG</code>	Execution of the calling process is not suspended if status is not immediately available for any child process.	<code>WUNTRACED</code>	The status of any child processes that are stopped, and whose status has not yet been reported since they stopped, are also reported to the requesting process.
<code>WNOHANG</code>	Execution of the calling process is not suspended if status is not immediately available for any child process.				
<code>WUNTRACED</code>	The status of any child processes that are stopped, and whose status has not yet been reported since they stopped, are also reported to the requesting process.				

wait3(3C)

When the `WNOHANG` option is specified and no processes have status to report, `wait3()` returns 0. The `WNOHANG` and `WUNTRACED` options may be combined by the bitwise OR operation of the two values.

The `wait4()` function is an extended interface. With a `pid` argument of 0, it is equivalent to `wait3()`. If `pid` has a nonzero value, then `wait4()` returns status only for the indicated process ID, but not for any other child processes. The status can be evaluated using the macros defined by `wait.h(3HEAD)`.

RETURN VALUES

If `wait3()` or `wait4()` returns due to a stopped or 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.

If `wait3()` or `wait4()` return due to the delivery of a signal to the calling process, `-1` is returned and `errno` is set to `EINTR`. If `WNOHANG` was 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`, 0 is returned. Otherwise, `-1` is returned and `errno` is set to indicate the error.

The `wait3()` and `wait4()` functions return 0 if `WNOHANG` is specified and there are no stopped or exited children, and return the process ID of the child process if they return due to a stopped or terminated child process. Otherwise, they return `-1` and set `errno` to indicate the error.

ERRORS

The `wait3()` and `wait4()` functions will fail and return immediately if:

<code>ECHILD</code>	The calling process has no existing unwaited-for child processes.
<code>EFAULT</code>	The <code>statusp</code> or <code>rusage</code> arguments point to an illegal address.
<code>EINTR</code>	The function was interrupted by a signal. The value of the location pointed to by <code>statusp</code> is undefined.
<code>EINVAL</code>	The value of <code>options</code> is not valid.

The `wait4()` function may fail if:

<code>ECHILD</code>	The process specified by <code>pid</code> does not exist or is not a child of the calling process.
---------------------	--

The `wait3()` and `wait4()` functions will terminate prematurely, return `-1`, and set `errno` to `EINTR` upon the arrival of a signal whose `SA_RESTART` bit in its flags field is not set (see `sigaction(2)`).

ATTRIBUTES

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

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

wait3(3C)

SEE ALSO kill(1), exit(2), waitid(2), waitpid(3C), getrusage(3C), signal(3C), signal.h(3HEAD), wait(3C), wait.h(3HEAD), proc(4)

NOTES If a parent process terminates without waiting on its children, the initialization process (process ID = 1) inherits the children.

The wait3() and wait4() functions are automatically restarted when a process receives a signal while awaiting termination of a child process, unless the SA_RESTART bit is not set in the flags for that signal.

wait(3C)

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 <code>wait()</code> 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 <code>wait()</code>, <code>waitpid(3C)</code>, or <code>waitid(2)</code> 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 <code>wait()</code>, return will be immediate.</p> <p>If <code>wait()</code> 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 <code>stat_loc</code>, the status of the child process is stored in the location pointed to by <code>stat_loc</code>. That status can be evaluated with the macros described on the <code>wait.h(3HEAD)</code> manual page.</p> <p>In the following, <code>status</code> is the object pointed to by <code>stat_loc</code>:</p> <ul style="list-style-type: none">■ If the child process terminated due to an <code>_exit()</code> call, the low order 8 bits of <code>status</code> will be 0 and the high order 8 bits will contain the low order 7 bits of the argument that the child process passed to <code>_exit()</code>; see <code>exit(2)</code>.■ If the child process terminated due to a signal, the high order 8 bits of <code>status</code> will be 0 and the low order 7 bits will contain the number of the signal that caused the termination. In addition, if <code>WCOREFLG</code> is set, a “core image” will have been produced; see <code>signal.h(3HEAD)</code> and <code>wait.h(3HEAD)</code>. <p>One instance of a <code>SIGCHLD</code> signal is queued for each child process whose status has changed. If <code>wait()</code> returns because the status of a child process is available, any pending <code>SIGCHLD</code> signal associated with the process ID of that child process is discarded. Any other pending <code>SIGCHLD</code> signals remain pending.</p> <p>If the calling process has <code>SA_NOCLDWAIT</code> set or has <code>SIGCHLD</code> set to <code>SIG_IGN</code>, and the process has no unwaited children that were transformed into zombie processes, it will block until all of its children terminate, and <code>wait()</code> will fail and set <code>errno</code> to <code>ECHILD</code>.</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 <code>intro(2)</code>.</p>		
RETURN VALUES	When <code>wait()</code> returns due to a terminated child process, the process ID of the child is returned to the calling process. Otherwise, <code>-1</code> is returned and <code>errno</code> is set to indicate the error.		
ERRORS	The <code>wait()</code> function will fail if: <table><tr><td><code>ECHILD</code></td><td>The calling process has no existing unwaited-for child processes.</td></tr></table>	<code>ECHILD</code>	The calling process has no existing unwaited-for child processes.
<code>ECHILD</code>	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 wanting to see the return results of such a call should use `waitpid(3C)` or `waitid(2)` instead of `wait()`. The `wait()` function is implemented as a call to `waitpid(-1, stat_loc, 0)`.

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

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

SEE ALSO `intro(2)`, `exec(2)`, `exit(2)`, `fork(2)`, `pause(2)`, `waitid(2)`, `ptrace(3C)`, `signal(3C)`, `signal.h(3HEAD)`, `waitpid(3C)`, `wait.h(3HEAD)`, `attributes(5)`

wait(3UCB)

NAME	wait, wait3, wait4, waitpid, WIFSTOPPED, WIFSIGNALED, WIFEXITED – wait for process to terminate or stop
SYNOPSIS	<pre><code>/usr/ucb/cc [flag ...] file ... #include <sys/wait.h> int wait(statusp); int *statusp; int waitpid(pid, statusp, options); int pid; int *statusp; int options; #include <sys/time.h> #include <sys/resource.h> int wait3(statusp, options, rusage); int *statusp; int options; struct rusage *rusage; int wait4(pid, statusp, options, rusage); int pid; int *statusp; int options; struct rusage *rusage; WIFSTOPPED(status); int status; WIFSIGNALED(status); int status; WIFEXITED(status); int status;</code></pre>
DESCRIPTION	<p>The <code>wait()</code> function delays its caller until a signal is received or one of its child processes terminates or stops due to tracing. If any child process has died or stopped due to tracing and this has not been reported using <code>wait()</code>, return is immediate, returning the process ID and exit status of one of those children. If that child process has died, it is discarded. If there are no children, return is immediate with the value <code>-1</code> returned. If there are only running or stopped but reported children, the calling process is blocked.</p> <p>If <code>status</code> is not a NULL pointer, then on return from a successful <code>wait()</code> call the status of the child process whose process ID is the return value of <code>wait()</code> is stored in the <code>wait()</code> union pointed to by <code>status</code>. The <code>w_status</code> member of that union is an <code>int</code>; it indicates the cause of termination and other information about the terminated process in the following manner:</p>

- If the low-order 8 bits of `w_status` are equal to 0177, the child process has stopped; the 8 bits higher up from the low-order 8 bits of `w_status` contain the number of the signal that caused the process to stop. See `ptrace(3C)` and `sigvec(3UCB)`.
- If the low-order 8 bits of `w_status` are non-zero and are not equal to 0177, the child process terminated due to a signal; the low-order 7 bits of `w_status` contain the number of the signal that terminated the process. In addition, if the low-order seventh bit of `w_status` (that is, bit 0200) is set, a “core image” of the process was produced; see `sigvec(3UCB)`.
- Otherwise, the child process terminated due to an `exit()` call; the 8 bits higher up from the low-order 8 bits of `w_status` contain the low-order 8 bits of the argument that the child process passed to `exit()`; see `exit(2)`.

`waitpid()` behaves identically to `wait()` if `pid` has a value of `-1` and `options` has a value of zero. Otherwise, the behavior of `waitpid()` is modified by the values of `pid` and `options` as follows:

`pid` specifies a set of child processes for which status is requested. `waitpid()` only returns the status of a child process from this set.

- If `pid` is equal to `-1`, status is requested for any child process. In this respect, `waitpid()` is then equivalent to `wait()`.
- If `pid` is greater than zero, it specifies the process ID of a single child process for which status is requested.
- If `pid` is equal to zero, status is requested for any child process whose process group ID is equal to that of the calling process.
- If `pid` is less than `-1`, status is requested for any child process whose process group ID is equal to the absolute value of `pid`.

`options` is constructed from the bitwise inclusive OR of zero or more of the following flags, defined in the header `<sys/wait.h>`:

<code>WNOHANG</code>	<code>waitpid()</code> does not suspend execution of the calling process if status is not immediately available for one of the child processes specified by <code>pid</code> .
<code>WUNTRACED</code>	The status of any child processes specified by <code>pid</code> that are stopped, and whose status has not yet been reported since they stopped, are also reported to the requesting process.

`wait3()` is an alternate interface that allows both non-blocking status collection and the collection of the status of children stopped by any means. The `status` parameter is defined as above. The `options` parameter is used to indicate the call should not block if there are no processes that have status to report (`WNOHANG`), and/or that children of the current process that are stopped due to a `SIGTTIN`, `SIGTTOU`, `SIGTSTP`, or `SIGSTOP` signal are eligible to have their status reported as well (`WUNTRACED`). A terminated child is discarded after it reports status, and a stopped process will not report its status more than once. If `rusage` is not a `NULL` pointer, a summary of the

wait(3UCB)

resources used by the terminated process and all its children is returned. Only the user time used and the system time used are currently available. They are returned in `rusage.ru_utime` and `rusage.ru_stime`, respectively.

When the `WNOHANG` option is specified and no processes have status to report, `wait3()` returns 0. The `WNOHANG` and `WUNTRACED` options may be combined by ORing the two values.

`wait4()` is another alternate interface. With a *pid* argument of 0, it is equivalent to `wait3()`. If *pid* has a nonzero value, then `wait4()` returns status only for the indicated process ID, but not for any other child processes.

`WIFSTOPPED`, `WIFSIGNALED`, `WIFEXITED`, are macros that take an argument *status*, of type `int`, as returned by `wait()`, or `wait3()`, or `wait4()`. `WIFSTOPPED` evaluates to true (1) when the process for which the `wait()` call was made is stopped, or to false (0) otherwise. `WIFSIGNALED` evaluates to true when the process was terminated with a signal. `WIFEXITED` evaluates to true when the process exited by using an `exit(2)` call.

RETURN VALUES

If `wait()` or `waitpid()` returns due to a stopped or terminated child process, the process ID of the child is returned to the calling process. Otherwise, a value of `-1` is returned and `errno` is set to indicate the error.

If `wait()` or `waitpid()` return due to the delivery of a signal to the calling process, a value of `-1` is returned and `errno` is set to `EINTR`. If `waitpid()` function was invoked with `WNOHANG` set in *options*, it has at least one child process specified by *pid* for which status is not available, and status is not available for any process specified by *pid*, a value of zero is returned. Otherwise, a value of `-1` is returned, and `errno` is set to indicate the error.

`wait3()` and `wait4()` returns 0 if `WNOHANG` is specified and there are no stopped or exited children, and returns the process ID of the child process if it returns due to a stopped or terminated child process. Otherwise, they returns a value of `-1` and sets `errno` to indicate the error.

ERRORS

`wait()`, `wait3()` or `wait4()` will fail and return immediately if one or more of the following are true:

`ECHILD` The calling process has no existing unwaited-for child processes.

`EFAULT` The *status* or *rusage* arguments point to an illegal address.

`waitpid()` may set `errno` to:

`ECHILD` The process or process group specified by *pid* does not exist or is not a child of the calling process.

`EINTR` The function was interrupted by a signal. The value of the location pointed to by *statusp* is undefined.

`EINVAL` The value of *options* is not valid.

wait(3UCB)

wait(), and wait3(), and wait4() will terminate prematurely, return -1, and set errno to EINTR upon the arrival of a signal whose SV_INTERRUPT bit in its flags field is set (see sigvec(3UCB) and siginterrupt(3UCB)). signal(3UCB), sets this bit for any signal it catches.

SEE ALSO exit(2), getrusage(3C), ptrace(3C), siginterrupt(3UCB), signal(3C), signal(3UCB), sigvec(3UCB), wait(3C), waitpid(3C)

NOTES Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multi-thread applications is unsupported.

If a parent process terminates without waiting on its children, the initialization process (process ID = 1) inherits the children.

The wait(), wait3(), and wait4() functions are automatically restarted when a process receives a signal while awaiting termination of a child process, unless the SV_INTERRUPT bit is set in the flags for that signal.

Calls to wait() with an argument of 0 should be cast to type 'int *', as in:

```
wait((int *)0)
```

Previous SunOS releases used union wait*statusp and union wait status in place of int *statusp and int status. The union contained a member w_status that could be treated in the same way as status.

Other members of the wait union could be used to extract this information more conveniently:

- If the w_stopval member had the value WSTOPPED, the child process had stopped; the value of the w_stopsig member was the signal that stopped the process.
- If the w_termsig member was non-zero, the child process terminated due to a signal; the value of the w_termsig member was the number of the signal that terminated the process. If the w_coredump member was non-zero, a core dump was produced.
- Otherwise, the child process terminated due to a call to exit(). The value of the w_retcode member was the low-order 8 bits of the argument that the child process passed to exit().

union wait is obsolete in light of the new specifications provided by *IEEE Std 1003.1-1988* and endorsed by *SVID89* and *XPG3*. SunOS Release 4.1 supports unionwait for backward compatibility, but it will disappear in a future release.

waitpid(3C)

NAME	waitpid – wait for child process to change state		
SYNOPSIS	<pre>#include <sys/types.h> #include <sys/wait.h> pid_t waitpid(pid_t <i>pid</i>, int *<i>stat_loc</i>, int <i>options</i>);</pre>		
DESCRIPTION	<p>The <code>waitpid()</code> 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 <code>waitpid()</code>, <code>wait(3C)</code>, or <code>waitid(2)</code> 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 <code>waitpid()</code>, 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 <code>(pid_t)-1</code>, status is requested for any child process. If <i>pid</i> is greater than <code>(pid_t)0</code>, it specifies the process ID of the child process for which status is requested.■ If <i>pid</i> is equal to <code>(pid_t)0</code> 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 <code>(pid_t)-1</code>, status is requested for any child process whose process group ID is equal to the absolute value of <i>pid</i>. <p>One instance of a SIGCHLD signal is queued for each child process whose status has changed. If <code>waitpid()</code> returns because the status of a child process is available, and WNOWAIT was not specified in <i>options</i>, any pending SIGCHLD signal associated with the process ID of that child process is discarded. Any other pending SIGCHLD signals remain pending.</p> <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 <code>waitpid()</code> will fail and set <code>errno</code> to ECHILD.</p> <p>If <code>waitpid()</code> returns because the status of a child process is available, then that status may be evaluated with the macros defined by <code>wait.h(3HEAD)</code>. 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> <table><tr><td>WCONTINUED</td><td>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.</td></tr></table>	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.
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.		

WNOHANG	The <code>waitpid()</code> function 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> .
WNOWAIT	Keep the process whose status is returned in <i>stat_loc</i> in a waitable state. The process may be waited for again with identical results.
WUNTRACED	The status of any child processes specified by <i>pid</i> that are stopped, and whose status has not yet been reported since they stopped, is also reported to the calling process. <code>WSTOPPED</code> is a synonym for <code>WUNTRACED</code> .

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 <i>pid</i> does not exist or is not a child of the calling process or can never be in the states specified by <i>options</i> .
EINTR	The <code>waitpid()</code> function was interrupted due to the receipt of a signal sent by the calling process.
EINVAL	An invalid value was specified for <i>options</i> .

USAGE With *options* equal to `0` and *pid* equal to `(pid_t)-1`, `waitpid()` is identical to `wait(3C)`. The `waitpid()` function is implemented as a call to the more general `waitid(2)` function.

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

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

SEE ALSO `intro(2)`, `exec(2)`, `exit(2)`, `fork(2)`, `pause(2)`, `sigaction(2)`, `ptrace(3C)`, `signal(3C)`, `siginfo.h(3HEAD)`, `wait(3C)`, `waitpid(3C)`, `wait.h(3HEAD)`, `attributes(5)`, `standards(5)`

walkcontext(3C)

NAME	walkcontext, printstack – walk stack pointed to by ucontext
SYNOPSIS	<pre>#include <ucontext.h> int walkcontext(const ucontext_t *uptr, int (*operate_func)(uintptr_t, int, void *), void *usrarg); int printstack(int fd);</pre>
DESCRIPTION	<p>The <code>walkcontext()</code> function walks the call stack pointed to by <code>uptr</code>, which can be obtained by a call to <code>getcontext(2)</code> or from a signal handler installed with the <code>SA_SIGINFO</code> flag. The <code>walkcontext()</code> function calls the user-supplied function <code>operate_func</code> for each routine found on the call stack and each signal handler invoked. The user function is passed three arguments: the PC at which the call or signal occurred, the signal number that occurred at this PC (0 if no signal occurred), and the third argument passed to <code>walkcontext()</code>. If the user function returns a non-zero value, <code>walkcontext()</code> returns without completing the callstack walk.</p> <p>The <code>printstack()</code> function uses <code>walkcontext()</code> to print a symbolic stack trace to the specified file descriptor. This is useful for reporting errors from signal handlers. The <code>printstack()</code> function uses <code>dladdr1()</code> (see dladdr(3C)) to obtain symbolic symbol names. As a result, only global symbols are reported as symbol names by <code>printstack()</code>.</p>
RETURN VALUES	Upon successful completion, <code>walkcontext()</code> and <code>printstack()</code> return 0. If <code>walkcontext()</code> cannot read the stack or the stack trace appears corrupted, both functions return -1.
ERRORS	No error values are defined.
USAGE	<p>The <code>walkcontext()</code> function is typically used to obtain information about the call stack for error reporting, performance analysis, or diagnostic purposes. Many library functions are not Async-Signal-Safe and should not be used from a signal handler. If <code>walkcontext()</code> is to be called from a signal handler, careful programming is required. In particular, stdio(3C) and malloc(3C) cannot be used.</p> <p>The <code>printstack()</code> function is Async-Signal-Safe and can be called from a signal handler. The output format from <code>printstack()</code> is unstable, as it varies with the scope of the routines.</p> <p>Tail-call optimizations on SPARC eliminate stack frames that would otherwise be present. For example, if the code is of the form</p> <pre>#include <stdio.h> main() { bar(); exit(0); } bar() {</pre>

```

        int a;
        a = foo(fileno(stdout));
        return (a);
    }

foo(int file)
{
    printstack(file);
}

```

compiling without optimization will yield a stack trace of the form

```

/tmp/q:foo+0x8
/tmp/q:bar+0x14
/tmp/q:main+0x4
/tmp/q:_start+0xb8

```

whereas with higher levels of optimization the output is

```

/tmp/q:main+0x10
/tmp/q:_start+0xb8

```

since both the call to `foo()` in `main` and the call to `bar()` in `foo()` are handled as tail calls that perform a return or restore in the delay slot. For further information, see *The SPARC Architecture Manual*.

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

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

SEE ALSO `intro(2)`, `getcontext(2)`, `sigaction(2)`, `dladdr(3C)`, `siginfo.h(3HEAD)`, `attributes(5)`

Weaver, David L. and Tom Germond, eds. *The SPARC Architecture Manual, Version 9*. Santa Clara: Prentice Hall, 2000.

watchmalloc(3MALLOC)

NAME	watchmalloc, cfree, memalign, valloc – debugging memory allocator
SYNOPSIS	<pre>#include <stdlib.h> void *malloc(size_t size); void free(void *ptr); void *realloc(void *ptr, size_t size); void *memalign(size_t alignment, size_t size); void *valloc(size_t size); void *calloc(size_t nelem, size_t elsize); void cfree(void *ptr, size_t nelem, size_t elsize); #include <malloc.h> int mallopt(int cmd, int value); struct mallinfo mallinfo(void);</pre>
DESCRIPTION	<p>The collection of <code>malloc()</code> functions in this shared object are an optional replacement for the standard versions of the same functions in the system C library. See malloc(3C). They provide a more strict interface than the standard versions and enable enforcement of the interface through the watchpoint facility of <code>/proc</code>. See <code>proc(4)</code>.</p> <p>Any dynamically linked application can be run with these functions in place of the standard functions if the following string is present in the environment (see <code>ld.so.1(1)</code>):</p> <pre>LD_PRELOAD=watchmalloc.so.1</pre> <p>The individual function interfaces are identical to the standard ones as described in malloc(3C). However, laxities provided in the standard versions are not permitted when the watchpoint facility is enabled (see <code>WATCHPOINTS</code> below):</p> <ul style="list-style-type: none">■ Memory may not be freed more than once.■ A pointer to freed memory may not be used in a call to <code>realloc()</code>.■ A call to <code>malloc()</code> immediately following a call to <code>free()</code> will not return the same space.■ Any reference to memory that has been freed yields undefined results. <p>To enforce these restrictions partially, without great loss in speed as compared to the watchpoint facility described below, a freed block of memory is overwritten with the pattern <code>0xdeadbeef</code> before returning from <code>free()</code>. The <code>malloc()</code> function returns with the allocated memory filled with the pattern <code>0xbaddcafe</code> as a precaution against applications incorrectly expecting to receive back unmodified memory from the last <code>free()</code>. The <code>calloc()</code> function always returns with the memory zero-filled.</p>

WATCHPOINTS

Entry points for `mallopt()` and `mallinfo()` are provided as empty routines, and are present only because some `malloc()` implementations provide them.

The watchpoint facility of `/proc` can be applied by a process to itself. The functions in `watchmalloc.so.1` use this feature if the following string is present in the environment:

```
MALLOC_DEBUG=WATCH
```

This causes every block of freed memory to be covered with `WA_WRITE` watched areas. If the application attempts to write any part of freed memory, it will trigger a watchpoint trap, resulting in a `SIGTRAP` signal, which normally produces an application core dump.

A header is maintained before each block of allocated memory. Each header is covered with a watched area, thereby providing a red zone before and after each block of allocated memory (the header for the subsequent memory block serves as the trailing red zone for its preceding memory block). Writing just before or just after a memory block returned by `malloc()` will trigger a watchpoint trap.

Watchpoints incur a large performance penalty. Requesting `MALLOC_DEBUG=WATCH` can cause the application to run 10 to 100 times slower, depending on the use made of allocated memory.

Further options are enabled by specifying a comma-separated string of options:

```
MALLOC_DEBUG=WATCH, RW, STOP
```

<code>WATCH</code>	Enables <code>WA_WRITE</code> watched areas as described above.
<code>RW</code>	Enables both <code>WA_READ</code> and <code>WA_WRITE</code> watched areas. An attempt either to read or write freed memory or the red zones will trigger a watchpoint trap. This incurs even more overhead and can cause the application to run up to 1000 times slower.
<code>STOP</code>	The process will stop showing a <code>FLTWATCH</code> machine fault if it triggers a watchpoint trap, rather than dumping core with a <code>SIGTRAP</code> signal. This allows a debugger to be attached to the live process at the point where it underwent the watchpoint trap. Also, the various <code>/proc</code> tools described in <code>proc(1)</code> can be used to examine the stopped process.

One of `WATCH` or `RW` must be specified, else the watchpoint facility is not engaged. `RW` overrides `WATCH`. Unrecognized options are silently ignored.

LIMITATIONS

Sizes of memory blocks allocated by `malloc()` are rounded up to the worst-case alignment size, 8 bytes for 32-bit processes and 16 bytes for 64-bit processes. Accessing the extra space allocated for a memory block is technically a memory violation but is in fact innocuous. Such accesses are not detected by the watchpoint facility of `watchmalloc`.

watchmalloc(3MALLOC)

Interposition of `watchmalloc.so.1` fails innocuously if the target application is statically linked with respect to its `malloc()` functions.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
MT-Level	MT-Safe

SEE ALSO `proc(1)`, `bsdmalloc(3MALLOC)`, `calloc(3C)`, `free(3C)`, `malloc(3C)`, `malloc(3MALLOC)`, `mapmalloc(3MALLOC)`, `memalign(3C)`, `realloc(3C)`, `valloc(3C)`, `libmapmalloc(3LIB)`, `proc(4)`, `attributes(5)`

NAME	wrtomb – convert a wide-character code to a character (restartable)				
SYNOPSIS	<pre>#include <stdio.h> size_t wrtomb(char *restrict s, wchar_t wc, mbstate_t *restrict ps);</pre>				
DESCRIPTION	<p>If <i>s</i> is a null pointer, the <code>wrtomb()</code> function is equivalent to the call:</p> <pre>wrtomb(buf, L'\0', ps)</pre> <p>where <i>buf</i> is an internal buffer.</p> <p>If <i>s</i> is not a null pointer, the <code>wrtomb()</code> function determines the number of bytes needed to represent the character that corresponds to the wide-character given by <i>wc</i> (including any shift sequences), and stores the resulting bytes in the array whose first element is pointed to by <i>s</i>. At most <code>MB_CUR_MAX</code> bytes are stored. If <i>wc</i> is a null wide-character, a null byte is stored, preceded by any shift sequence needed to restore the initial shift state. The resulting state described is the initial conversion state.</p> <p>If <i>ps</i> is a null pointer, the <code>wrtomb()</code> function uses its own internal <code>mbstate_t</code> object, which is initialized at program startup to the initial conversion state. Otherwise, the <code>mbstate_t</code> object pointed to by <i>ps</i> is used to completely describe the current conversion state of the associated character sequence. Solaris will behave as if no function defined in the Solaris Reference Manual calls <code>wrtomb()</code>.</p> <p>The behavior of this function is affected by the <code>LC_CTYPE</code> category of the current locale. See <code>environ(5)</code>.</p>				
RETURN VALUES	The <code>wrtomb()</code> function returns the number of bytes stored in the array object (including any shift sequences). When <i>wc</i> is not a valid wide-character, an encoding error occurs. In this case, the function stores the value of the macros <code>EILSEQ</code> in <code>errno</code> and returns <code>(size_t)-1</code> ; the conversion state is undefined.				
ERRORS	<p>The <code>wrtomb()</code> function may fail if:</p> <table border="0" style="margin-left: 20px;"> <tr> <td style="padding-right: 20px;"><code>EINVAL</code></td> <td>The <i>ps</i> argument points to an object that contains an invalid conversion state.</td> </tr> <tr> <td><code>EILSEQ</code></td> <td>Invalid wide-character code is detected.</td> </tr> </table>	<code>EINVAL</code>	The <i>ps</i> argument points to an object that contains an invalid conversion state.	<code>EILSEQ</code>	Invalid wide-character code is detected.
<code>EINVAL</code>	The <i>ps</i> argument points to an object that contains an invalid conversion state.				
<code>EILSEQ</code>	Invalid wide-character code is detected.				
USAGE	If <i>ps</i> is not a null pointer, <code>wrtomb()</code> uses the <code>mbstate_t</code> object pointed to by <i>ps</i> and the function can be used safely in multithreaded applications, as long as <code>setlocale(3C)</code> is not being called to change the locale. If <i>ps</i> is a null pointer, <code>wrtomb()</code> uses its internal <code>mbstate_t</code> object and the function is Unsafe in multithreaded applications.				
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:				

wcrtomb(3C)

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	See NOTES below

SEE ALSO [mbsinit\(3C\)](#), [setlocale\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#), [environ\(5\)](#)

NAME	wscoll, wscoll – wide character string comparison using collating information								
SYNOPSIS	<pre>#include <wchar.h> int wscoll(const wchar_t *ws1, const wchar_t *ws2); int wscoll(const wchar_t *ws1, const wchar_t *ws2);</pre>								
DESCRIPTION	<p>The <code>wscoll()</code> and <code>wscoll()</code> functions compare the wide character string pointed to by <code>ws1</code> to the wide character string pointed to by <code>ws2</code>, both interpreted as appropriate to the <code>LC_COLLATE</code> category of the current locale.</p> <p>The <code>wscoll()</code> and <code>wscoll()</code> functions do not change the setting of <code>errno</code> if successful.</p> <p>An application wanting to check for error situations should set <code>errno</code> to 0 before calling <code>wscoll()</code> or <code>wscoll()</code>. If <code>errno</code> is non-zero on return, an error has occurred.</p>								
RETURN VALUES	<p>Upon successful completion, <code>wscoll()</code> and <code>wscoll()</code> return an integer greater than, equal to, or less than 0, depending upon whether the wide character string pointed to by <code>ws1</code> is greater than, equal to, or less than the wide character string pointed to by <code>ws2</code>, when both are interpreted as appropriate to the current locale. On error, <code>wscoll()</code> and <code>wscoll()</code> may set <code>errno</code>, but no return value is reserved to indicate an error.</p>								
ERRORS	<p>The <code>wscoll()</code> and <code>wscoll()</code> functions may fail if:</p> <p><code>EINVAL</code> The <code>ws1</code> or <code>ws2</code> arguments contain wide character codes outside the domain of the collating sequence.</p>								
USAGE	The <code>wcsxfrm(3C)</code> and <code>wscmp(3C)</code> functions should be used for sorting large lists.								
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:								
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>CSI</td> <td>Enabled</td> </tr> <tr> <td>Interface Stability</td> <td><code>wscoll()</code> is Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	<code>wscoll()</code> is Standard	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
CSI	Enabled								
Interface Stability	<code>wscoll()</code> is Standard								
MT-Level	MT-Safe with exceptions								
SEE ALSO	<p>The <code>wscoll()</code> and <code>wscoll()</code> functions can be used safely in multithreaded applications as long as <code>setlocale(3C)</code> is not being called to change the locale.</p> <p><code>setlocale(3C)</code>, <code>wscmp(3C)</code>, <code>wcsxfrm(3C)</code>, <code>attributes(5)</code>, <code>standards(5)</code></p>								

wcsftime(3C)

NAME	wcsftime – convert date and time to wide character string								
SYNOPSIS	#include <wchar.h>								
XPG4 and SUS	size_t wcsftime (wchar_t *wcs, size_t maxsize, const char *format, const struct tm *timptr);								
Default and other standards	size_t wcsftime (wchar_t *restrict wcs, size_t maxsize, const wchar_t *restrict format, const struct tm *restrict timptr);								
DESCRIPTION	<p>The <code>wcsftime()</code> function is equivalent to the strftime(3C) function, except that:</p> <ul style="list-style-type: none">■ The argument <code>wcs</code> points to the initial element of an array of wide-characters into which the generated output is to be placed.■ The argument <code>maxsize</code> indicates the maximum number of wide-characters to be placed in the output array.■ The argument <code>format</code> is a wide-character string and the conversion specifications are replaced by corresponding sequences of wide-characters.■ The return value indicates the number of wide-characters placed in the output array. <p>If copying takes place between objects that overlap, the behavior is undefined.</p>								
RETURN VALUES	<p>If the total number of resulting wide character codes (including the terminating null wide-character code) is no more than <code>maxsize</code>, <code>wcsftime()</code> returns the number of wide-character codes placed into the array pointed to by <code>wcs</code>, not including the terminating null wide-character code. Otherwise, 0 is returned and the contents of the array are indeterminate.</p> <p>The <code>wcsftime()</code> function uses malloc(3C) and should <code>malloc()</code> fail, <code>errno</code> will be set by <code>malloc()</code>.</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>CSI</td><td>Enabled</td></tr><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe with exceptions</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	Standard	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
CSI	Enabled								
Interface Stability	Standard								
MT-Level	MT-Safe with exceptions								
SEE ALSO	malloc(3C) , setlocale(3C) , strftime(3C) , attributes(5) , standards(5)								
NOTES	The <code>wcsftime()</code> function can be used safely in multithreaded applications, as long as setlocale(3C) is not being called to change the locale.								

NAME	wcsrtombs – convert a wide-character string to a character string (restartable)				
SYNOPSIS	<pre>#include <wchar.h> size_t wcsrtombs(char *restrict <i>dst</i>, const wchar_t **restrict <i>src</i>, size_t <i>len</i>, mbstate_t *restrict <i>ps</i>);</pre>				
DESCRIPTION	<p>The <code>wcsrtombs()</code> function converts a sequence of wide-characters from the array indirectly pointed to by <i>src</i> into a sequence of corresponding characters, beginning in the conversion state described by the object pointed to by <i>ps</i>. If <i>dst</i> is not a null pointer, the converted characters are then stored into the array pointed to by <i>dst</i>. Conversion continues up to and including a terminating null wide-character, which is also stored. Conversion stops earlier in the following cases:</p> <ul style="list-style-type: none"> ■ When a code is reached that does not correspond to a valid character. ■ When the next character would exceed the limit of <i>len</i> total bytes to be stored in the array pointed to by <i>dst</i> (and <i>dst</i> is not a null pointer). <p>Each conversion takes place as if by a call to the <code>wcrtomb()</code> function.</p> <p>If <i>dst</i> is not a null pointer, the pointer object pointed to by <i>src</i> is assigned either a null pointer (if conversion stopped due to reaching a terminating null wide-character) or the address just past the last wide-character converted (if any). If conversion stopped due to reaching a terminating null wide-character, the resulting state described is the initial conversion state.</p> <p>If <i>ps</i> is a null pointer, the <code>wcsrtombs()</code> function uses its own internal <code>mbstate_t</code> object, which is initialized at program startup to the initial conversion state. Otherwise, the <code>mbstate_t</code> object pointed to by <i>ps</i> is used to completely describe the current conversion state of the associated character sequence. Solaris will behave as if no function defined in the Solaris Reference Manual calls <code>wcsrtombs()</code>.</p> <p>The behavior of this function is affected by the <code>LC_CTYPE</code> category of the current locale. See <code>environ(5)</code>.</p>				
RETURN VALUES	<p>If conversion stops because a code is reached that does not correspond to a valid character, an encoding error occurs. In this case, the <code>wcsrtombs()</code> function stores the value of the macro <code>EILSEQ</code> in <code>errno</code> and returns <code>(size_t)-1</code>; the conversion state is undefined. Otherwise, it returns the number of bytes in the resulting character sequence, not including the terminating null (if any).</p>				
ERRORS	<p>The <code>wcsrtombs()</code> function may fail if:</p> <table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;"><code>EINVAL</code></td> <td>The <i>ps</i> argument points to an object that contains an invalid conversion state.</td> </tr> <tr> <td><code>EILSEQ</code></td> <td>A wide-character code does not correspond to a valid character.</td> </tr> </table>	<code>EINVAL</code>	The <i>ps</i> argument points to an object that contains an invalid conversion state.	<code>EILSEQ</code>	A wide-character code does not correspond to a valid character.
<code>EINVAL</code>	The <i>ps</i> argument points to an object that contains an invalid conversion state.				
<code>EILSEQ</code>	A wide-character code does not correspond to a valid character.				
USAGE	<p>If <i>ps</i> is not a null pointer, <code>wcsrtombs()</code> uses the <code>mbstate_t</code> object pointed to by <i>ps</i> and the function can be used safely in multithreaded applications, as long as <code>setlocale(3C)</code> is not being called to change the locale. If <i>ps</i> is a null pointer, <code>wcsrtombs()</code> uses its internal <code>mbstate_t</code> object and the function is Unsafe in</p>				

wcsrtombs(3C)

multithreaded applications.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	See NOTES below

SEE ALSO `mbsinit(3C)`, `setlocale(3C)`, `wcrtomb(3C)`, `attributes(5)`, `environ(5)`, `standards(5)`

NAME	wcsstr – find a wide-character substring
SYNOPSIS	<pre>#include <wchar.h> wchar_t *wcsstr(const wchar_t *restrict ws1, const wchar_t *restrict ws2);</pre>
ISO C++	<pre>#include <wchar.h> const wchar_t *wcsstr(const wchar_t *ws1, const wchar_t *ws2); #include <cwchar> wchar_t *std::wcsstr(wchar_t *ws1, const wchar_t *ws2);</pre>
DESCRIPTION	The <code>wcsstr()</code> function locates the first occurrence in the wide-character string pointed to by <code>ws1</code> of the sequence of wide-characters (excluding the terminating null wide-character) in the wide-character string pointed to by <code>ws2</code> .
RETURN VALUES	On successful completion, <code>wcsstr()</code> returns a pointer to the located wide-character string, or a null pointer if the wide-character string is not found. If <code>ws2</code> points to a wide-character string with zero length, the function returns <code>ws1</code> .
ERRORS	No errors are defined.
ATTRIBUTES	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO [wchr\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

wcstod(3C)

NAME	<code>wcstod</code> , <code>wcstof</code> , <code>wcstold</code> , <code>wstod</code> , <code>watof</code> – convert wide character string to floating-point number
SYNOPSIS	<pre>#include <wchar.h> double wcstod(const wchar_t *restrict <i>nptr</i>, wchar_t **restrict <i>endptr</i>) ; float wcstof(const wchar_t *restrict <i>nptr</i>, wchar_t **restrict <i>endptr</i>) ; long double wcstold(const wchar_t *restrict <i>nptr</i>, wchar_t **restrict <i>endptr</i>) ; double wstod(const wchar_t *<i>nptr</i>, wchar_t **<i>endptr</i>) ; double watof(wchar_t *<i>nptr</i>) ;</pre>
DESCRIPTION	<p>The <code>wcstod()</code>, <code>wcstof()</code>, and <code>wcstold()</code> functions convert the initial portion of the wide-character string pointed to by <i>nptr</i> to double, float, and long double representation, respectively. They first decompose the input wide-character string into three parts:</p> <ol style="list-style-type: none">1. An initial, possibly empty, sequence of white-space wide-character codes (as specified by iswspace(3C))2. A subject sequence interpreted as a floating-point constant or representing infinity or NaN3. A final wide-character string of one or more unrecognized wide-character codes, including the terminating null wide-character code of the input wide-character string. <p>Then they attempt to convert the subject sequence to a floating-point number, and return the result.</p> <p>The expected form of the subject sequence is an optional plus or minus sign, then one of the following:</p> <ul style="list-style-type: none">■ A non-empty sequence of decimal digits optionally containing a radix character, then an optional exponent part■ A 0x or 0X, then a non-empty sequence of hexadecimal digits optionally containing a radix character, then an optional binary exponent part■ One of INF or INFINITY, or any other wide string equivalent except for case■ One of NAN or NAN(<i>n-wchar-sequence_{opt}</i>), or any other wide string ignoring case in the NAN part, where: <pre>n-wchar-sequence: digit nondigit n-wchar-sequence digit n-wchar-sequence nondigit</pre>

In default mode for `wcstod()`, only decimal, INF/INFINITY, and NAN/NAN(*n-char-sequence*) forms are recognized. In C99/SUSv3 mode, hexadecimal strings are also recognized.

In default mode for `wcstod()`, the *n-char-sequence* in the NAN(*n-char-sequence*) form can contain any character except ')' (right parenthesis) or '\0' (null). In C99/SUSv3 mode, the *n-char-sequence* can contain only upper and lower case letters, digits, and '_' (underscore).

The `wcstof()` and `wcstold()` functions always function in C99/SUSv3-conformant mode.

The subject sequence is defined as the longest initial subsequence of the input wide string, starting with the first non-white-space wide character, that is of the expected form. The subject sequence contains no wide characters if the input wide string is not of the expected form.

If the subject sequence has the expected form for a floating-point number, the sequence of wide characters starting with the first digit or the radix character (whichever occurs first) is interpreted as a floating constant according to the rules of the C language, except that the radix character is used in place of a period, and that if neither an exponent part nor a radix character appears in a decimal floating-point number, or if a binary exponent part does not appear in a hexadecimal floating-point number, an exponent part of the appropriate type with value zero is assumed to follow the last digit in the string. If the subject sequence begins with a minus sign, the sequence is interpreted as negated. A wide-character sequence INF or INFINITY is interpreted as an infinity. A wide-character sequence NAN or NAN(*n-wchar-sequence_{opt}*) is interpreted as a quiet NaN. A pointer to the final wide string is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

If the subject sequence has either the decimal or hexadecimal form, the value resulting from the conversion is rounded correctly according to the prevailing floating point rounding direction mode. The conversion also raises floating point inexact, underflow, or overflow exceptions as appropriate.

The radix character is defined in the program's locale (category LC_NUMERIC). In the POSIX locale, or in a locale where the radix character is not defined, the radix character defaults to a period ('.').

If the subject sequence is empty or does not have the expected form, no conversion is performed; the value of *nptr* is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

The `wcstod()` function does not change the setting of `errno` if successful.

The `wstod()` function is identical to `wcstod()`.

The `watof(str)` function is equivalent to `wstod(nptr, (wchar_t **)NULL)`.

RETURN VALUES

Upon successful completion, these functions return the converted value. If no conversion could be performed, 0 is returned.

wcstod(3C)

If the correct value is outside the range of representable values, `±HUGE_VAL`, `±HUGE_VALF`, or `±HUGE_VALL` is returned (according to the sign of the value), a floating point overflow exception is raised, and `errno` is set to `ERANGE`.

If the correct value would cause an underflow, the correctly rounded result (which may be normal, subnormal, or zero) is returned, a floating point underflow exception is raised, and `errno` is set to `ERANGE`.

ERRORS The `wcstod()` and `wctod()` functions will fail if:

`ERANGE` The value to be returned would cause overflow or underflow.

The `wcstod()` and `wctod()` functions may fail if:

`EINVAL` No conversion could be performed.

USAGE Because 0 is returned on error and is also a valid return on success, an application wishing to check for error situations should set `errno` to 0 call `wcstod()`, `wctod()`, `wcstold()`, or `wctold()`, then check `errno` and if it is non-zero, assume an error has occurred.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	<code>wcstod()</code> , <code>wctod()</code> , and <code>wcstold()</code> are Standard.
MT-Level	MT-Safe

SEE ALSO `iswspace(3C)`, `localeconv(3C)`, `scanf(3C)`, `setlocale(3C)`, `wcstol(3C)`, `attributes(5)`, `standards(5)`

NAME	wcstoimax, wcstoumax – convert wide-character string to integer type						
SYNOPSIS	<pre>#include <stddef.h> #include <inttypes.h> intmax_t wcstoimax(const wchar_t *restrict <i>nptr</i>, wchar_t **restrict <i>endptr</i>, int <i>base</i>); uintmax_t wcstoumax(const wchar_t *restrict <i>nptr</i>, wchar_t **restrict <i>endptr</i>, int <i>base</i>);</pre>						
DESCRIPTION	These functions are equivalent to the wcstol(3C) , wcstoll(3C) , wcstoul(3C) , and wcstoull(3C) functions, respectively, except that the initial portion of the wide string is converted to <code>intmax_t</code> and <code>uintmax_t</code> representation, respectively.						
RETURN VALUES	These functions return the converted value, if any. If no conversion could be performed, 0 is returned. If the correct value is outside the range of representable values, <code>{INTMAX_MAX}</code> , <code>{INTMAX_MIN}</code> , or <code>{UINTMAX_MAX}</code> is returned (according to the return type and sign of the value), and <code>errno</code> is set to <code>ERANGE</code> .						
ERRORS	<p>These functions will fail if:</p> <p><code>EINVAL</code> The value of <i>base</i> is not supported.</p> <p><code>ERANGE</code> The value to be returned is not representable.</p> <p>These functions may fail if:</p> <p><code>EINVAL</code> No conversion could be performed.</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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	wcstol(3C) , wcstoul(3C) , attributes(5) , standards(5)						

wcstol(3C)

NAME	<code>wcstol</code> , <code>wcstoll</code> , <code>wstol</code> , <code>watol</code> , <code>watoll</code> , <code>watoi</code> – convert wide character string to long integer
SYNOPSIS	<pre>#include <wchar.h> long wcstol(const wchar_t *restrict <i>nptr</i>, wchar_t **restrict <i>endptr</i>, int <i>base</i>); long long wcstoll(const wchar_t *restrict <i>nptr</i>, wchar_t **restrict <i>endptr</i>, int <i>base</i>); #include <wider.h> long wstol(const wchar_t *<i>nptr</i>, wchar_t **<i>endptr</i>, int <i>base</i>); long watol(wchar_t *<i>nptr</i>); long long watoll(wchar_t *<i>nptr</i>); int watoi(wchar_t *<i>nptr</i>);</pre>
DESCRIPTION	<p>The <code>wcstol()</code> and <code>wcstoll()</code> functions convert the initial portion of the wide character string pointed to by <i>nptr</i> to long and long long representation, respectively. They first decompose the input string into three parts:</p> <ol style="list-style-type: none">1. an initial, possibly empty, sequence of white-space wide-character codes (as specified by iswspace(3C))2. a subject sequence interpreted as an integer represented in some radix determined by the value of <i>base</i>3. a final wide character string of one or more unrecognised wide character codes, including the terminating null wide-character code of the input wide character string <p>They then attempt to convert the subject sequence to an integer, and return the result.</p> <p>If the value of <i>base</i> is 0, the expected form of the subject sequence is that of a decimal constant, octal constant or hexadecimal constant, any of which may be preceded by a '+' or '-' sign. A decimal constant begins with a non-zero digit, and consists of a sequence of decimal digits. An octal constant consists of the prefix '0' optionally followed by a sequence of the digits '0' to '7' only. A hexadecimal constant consists of the prefix '0x' or '0X' followed by a sequence of the decimal digits and letters 'a' (or 'A') to 'f' (or 'F') with values 10 to 15 respectively.</p> <p>If the value of <i>base</i> is between 2 and 36, the expected form of the subject sequence is a sequence of letters and digits representing an integer with the radix specified by <i>base</i>, optionally preceded by a '+' or '-' sign, but not including an integer suffix. The letters from 'a' (or 'A') to 'z' (or 'Z') inclusive are ascribed the values 10 to 35; only letters whose ascribed values are less than that of <i>base</i> are permitted. If the value of <i>base</i> is 16, the wide-character code representations of '0x' or '0X' may optionally precede the sequence of letters and digits, following the sign if present.</p>

The subject sequence is defined as the longest initial subsequence of the input wide character string, starting with the first non-white-space wide-character code, that is of the expected form. The subject sequence contains no wide-character codes if the input wide character string is empty or consists entirely of white-space wide-character code, or if the first non-white-space wide-character code is other than a sign or a permissible letter or digit.

If the subject sequence has the expected form and the value of *base* is 0, the sequence of wide-character codes starting with the first digit is interpreted as an integer constant. If the subject sequence has the expected form and the value of *base* is between 2 and 36, it is used as the base for conversion, ascribing to each letter its value as given above. If the subject sequence begins with a minus sign (-), the value resulting from the conversion is negated. A pointer to the final wide character string is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

If the subject sequence is empty or does not have the expected form, no conversion is performed; the value of *nptr* is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

These functions do not change the setting of `errno` if successful.

Since 0, `{LONG_MIN}` or `{LLONG_MIN}`, and `{LONG_MAX}` or `{LLONG_MAX}` are returned on error and are also valid returns on success, an application wanting to check for error situations should set `errno` to 0, call one of these functions, then check `errno`.

The `wstol()` function is equivalent to `wcstol()`.

The `watol()` function is equivalent to `wstol(str, (wchar_t **)NULL, 10)`.

The `watoll()` function is the long-long (double long) version of `watol()`.

The `watoi()` function is equivalent to `(int)watol()`.

RETURN VALUES

Upon successful completion, these functions return the converted value, if any. If no conversion could be performed, 0 is returned and `errno` may be set to indicate the error. If the correct value is outside the range of representable values, `{LONG_MIN}`, `{LONG_MAX}`, `{LLONG_MIN}`, or `{LLONG_MAX}` is returned (according to the sign of the value), and `errno` is set to `ERANGE`.

ERRORS

These functions will fail if:

<code>EINVAL</code>	The value of <i>base</i> is not supported.
<code>ERANGE</code>	The value to be returned is not representable.

These functions may fail if:

<code>EINVAL</code>	No conversion could be performed.
---------------------	-----------------------------------

wcstol(3C)

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	<code>wcstol()</code> and <code>wcstoll()</code> are Standard.
MT-Level	MT-Safe

SEE ALSO [iswalpha\(3C\)](#), [iswspace\(3C\)](#), [scanf\(3C\)](#), [wcstod\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NOTES Truncation from `long long` to `long` can take place upon assignment or by an explicit cast.

NAME	wcstombs – convert a wide-character string to a character string								
SYNOPSIS	<pre>#include <stdlib.h> size_t wcstombs(char *restrict s, const wchar_t *restrict pwcs, size_t n);</pre>								
DESCRIPTION	<p>The <code>wcstombs()</code> function converts the sequence of wide-character codes from the array pointed to by <code>pwcs</code> into a sequence of characters and stores these characters into the array pointed to by <code>s</code>, stopping if a character would exceed the limit of <code>n</code> total bytes or if a null byte is stored. Each wide-character code is converted as if by a call to <code>wctomb(3C)</code>.</p> <p>The behavior of this function is affected by the <code>LC_CTYPE</code> category of the current locale.</p> <p>No more than <code>n</code> bytes will be modified in the array pointed to by <code>s</code>. If copying takes place between objects that overlap, the behavior is undefined. If <code>s</code> is a null pointer, <code>wcstombs()</code> returns the length required to convert the entire array regardless of the value of <code>n</code>, but no values are stored.</p>								
RETURN VALUES	<p>If a wide-character code is encountered that does not correspond to a valid character (of one or more bytes each), <code>wcstombs()</code> returns <code>(size_t)-1</code>. Otherwise, <code>wcstombs()</code> returns the number of bytes stored in the character array, not including any terminating null byte. The array will not be null-terminated if the value returned is <code>n</code>.</p>								
ERRORS	<p>The <code>wcstombs()</code> function may fail if:</p> <p>EILSEQ A wide-character code does not correspond to a valid character.</p>								
ATTRIBUTES	<p>See <code>attributes(5)</code> 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>CSI</td> <td>Enabled</td> </tr> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
CSI	Enabled								
Interface Stability	Standard								
MT-Level	MT-Safe								
SEE ALSO	<p><code>mblen(3C)</code>, <code>mbstowcs(3C)</code>, <code>mbtowc(3C)</code>, <code>setlocale(3C)</code>, <code>wctomb(3C)</code>, <code>attributes(5)</code>, <code>standards(5)</code></p>								

wcstoul(3C)

NAME	wcstoul, wcstoull – convert wide-character string to unsigned long
SYNOPSIS	<pre>#include <wchar.h> unsigned long wcstoul(const wchar_t *restrict <i>nptr</i>, wchar_t **restrict <i>endptr</i>, int <i>base</i>); unsigned long long wcstoull(const wchar_t *restrict <i>nptr</i>, wchar_t **restrict <i>endptr</i>, int <i>base</i>);</pre>
DESCRIPTION	<p>The <code>wcstoul()</code> and <code>wcstoull()</code> functions convert the initial portion of the wide-character string pointed to by <i>nptr</i> to unsigned long and unsigned long long representation, respectively. First they decompose the input wide-character string into three parts:</p> <ol style="list-style-type: none">1. An initial, possibly empty, sequence of white-space wide-character codes (as specified by the function <code>iswspace(3C)</code>)2. A subject sequence interpreted as an integer represented in some radix determined by the value of <i>base</i>3. a final wide-character string of one or more unrecognized wide-character codes, including the terminating null wide-character code of the input wide character string <p>They then attempt to convert the subject sequence to an unsigned integer and return the result.</p> <p>If the value of <i>base</i> is 0, the expected form of the subject sequence is that of a decimal constant, an octal constant, or a hexadecimal constant, any of which may be preceded by a '+' or a '-' sign. A decimal constant begins with a non-zero digit, and consists of a sequence of decimal digits. An octal constant consists of the prefix '0', optionally followed by a sequence of the digits '0' to '7' only. A hexadecimal constant consists of the prefix '0x' or '0X', followed by a sequence of the decimal digits and letters 'a' (or 'A') to 'f' (or 'F'), with values 10 to 15, respectively.</p> <p>If the value of <i>base</i> is between 2 and 36, the expected form of the subject sequence is a sequence of letters and digits representing an integer with the radix specified by <i>base</i>, optionally preceded by a '+' or a '-' sign, but not including an integer suffix. The letters from 'a' (or 'A') to 'z' (or 'Z') inclusive are ascribed the values 10 to 35; only letters whose ascribed values are less than that of <i>base</i> are permitted. If the value of <i>base</i> is 16, the wide-character codes '0x' or '0X' may optionally precede the sequence of letters and digits, following the sign, if present.</p> <p>The subject sequence is defined as the longest initial subsequence of the input wide-character string, starting with the first wide-character code that is not a white space and is of the expected form. The subject sequence contains no wide-character codes if the input wide-character string is empty or consists entirely of white-space wide-character codes, or if the first wide-character code that is not a white space is other than a sign or a permissible letter or digit.</p>

If the subject sequence has the expected form and the value of *base* is 0, the sequence of wide-character codes starting with the first digit is interpreted as an integer constant. If the subject sequence has the expected form and the value of *base* is between 2 and 36, it is used as the base for conversion, ascribing to each letter its value as given above. If the subject sequence begins with a minus sign, the value resulting from the conversion is negated. A pointer to the final wide character string is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

If the subject sequence is empty or does not have the expected form, no conversion is performed; the value of *nptr* is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

The `wcstoul()` function does not change the setting of `errno` if successful.

Since 0, `{ULONG_MAX}`, and `{ULLONG_MAX}` are returned on error and 0 is also a valid return on success, an application wanting to check for error situations should set `errno` to 0, then call `wcstoul()` or `wcstoull()`, then check `errno`.

The `wcstoul()` and `wcstoull()` functions do not change the setting of `errno` if successful.

RETURN VALUE

Upon successful completion, `wcstoul()` and `wcstoull()` return the converted value, if any. If no conversion could be performed, 0 is returned and `errno` may be set to indicate the error. If the correct value is outside the range of representable values, `{ULONG_MAX}` or `{ULLONG_MAX}`, respectively, is returned and `errno` is set to `ERANGE`.

ERRORS

The `wcstoul()` and `wcstoull()` functions will fail if:

`EINVAL` The value of *base* is not supported.
`ERANGE` The value to be returned is not representable.

The `wcstoul()` and `wcstoull()` functions may fail if:

`EINVAL` No conversion could be performed.

USAGE

Unlike `wcstod(3C)` and `wcstol(3C)`, `wcstoul()` and `wcstoull()` must always return a non-negative number; using the return value of `wcstoul()` for out-of-range numbers with `wcstoul()` or `wcstoull()` could cause more severe problems than just loss of precision if those numbers can ever be negative.

ATTRIBUTES

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO

`isspace(3C)`, `iswalph(3C)`, `scanf(3C)`, `wcstod(3C)`, `wcstol(3C)`, `attributes(5)`, `standards(5)`

wcstring(3C)

NAME	wcstring, wcsat, wscat, wcsncat, wsncat, wscmp, wscmp, wcsncmp, wsncmp, wcsncpy, wcsncpy, wcsncpy, wsncpy, wcslen, wslen, wcschr, wschr, wcsrchr, wsrchr, windex, wrindex, wcpbrk, wcpbrk, wswcs, wcsspn, wssp, wcsspn, wcsspn, wstok, wstok – wide-character string operations
SYNOPSIS	<pre>#include <wchar.h> wchar_t *wscat(wchar_t *ws1, const wchar_t *ws2); wchar_t *wcsncat(wchar_t *restrict ws1, const wchar_t *restrict ws2, size_t n); int wscmp(const wchar_t *ws1, const wchar_t *ws2); int wcsncmp(const wchar_t *ws1, const wchar_t *ws2, size_t n); wchar_t *wcsncpy(wchar_t *ws1, const wchar_t *ws2); wchar_t *wcsncpy(wchar_t *restrict ws1, const wchar_t *restrict ws2, size_t n); size_t wcslen(const wchar_t *ws); wchar_t *wcschr(const wchar_t *ws, wchar_t wc); wchar_t *wcsrchr(const wchar_t *ws, wchar_t wc); wchar_t *wcpbrk(const wchar_t *ws1, const wchar_t *ws2); wchar_t *wswcs(const wchar_t *ws1, const wchar_t *ws2); size_t wcsspn(const wchar_t *ws1, const wchar_t *ws2); size_t wcsspn(const wchar_t *ws1, const wchar_t *ws2); wchar_t *wstok(wchar_t *restrict ws1, const wchar_t *restrict ws2);</pre>
XPG4, SUS, SUSv2, SUSv3	
Default and other standards	<pre>wchar_t *wstok(wchar_t *ws1, const wchar_t *ws2, wchar_t **ptr); #include <wchar.h> wchar_t *wscat(wchar_t *ws1, const wchar_t *ws2); wchar_t *wsncat(wchar_t *ws1, const wchar_t *ws2, size_t n); int wscmp(const wchar_t *ws1, const wchar_t *ws2); int wsncmp(const wchar_t *ws1, const wchar_t *ws2, size_t n); wchar_t *wscpy(wchar_t *ws1, const wchar_t *ws2); wchar_t *wsncpy(wchar_t *ws1, const wchar_t *ws2, size_t n); size_t wslen(const wchar_t *ws); wchar_t *wschr(const wchar_t *ws, wchar_t wc); wchar_t *wsrchr(const wchar_t *ws, wchar_t wc); wchar_t *wcpbrk(const wchar_t *ws1, const wchar_t *ws2);</pre>

	<pre> size_t wsspn(const wchar_t *ws1, const wchar_t *ws2); size_t wscspn(const wchar_t *ws1, const wchar_t *ws2); wchar_t *wstok(wchar_t *ws1, const wchar_t *ws2); wchar_t *windex(const wchar_t *ws, wchar_t wc); wchar_t *wrindex(const wchar_t *ws, wchar_t wc); </pre>
ISO C++	<pre> #include <wchar.h> const wchar_t *wcschr(const wchar_t *ws, wchar_t wc); const wchar_t *wcspbrk(const wchar_t *ws1, const wchar_t *ws2); const wchar_t *wcsrchr(const wchar_t *ws, wchar_t wc); #include <cwchar> wchar_t *std::wcschr(wchar_t *ws, wchar_t wc); wchar_t *std::wcspbrk(wchar_t *ws1, const wchar_t *ws2); wchar_t *std::wcsrchr(wchar_t *ws, wchar_t wc); </pre>
DESCRIPTION	<p>These functions operate on wide-character strings terminated by <code>wchar_t NULL</code> characters. During appending or copying, these routines do not check for an overflow condition of the receiving string. In the following, <i>ws</i>, <i>ws1</i>, and <i>ws2</i> point to wide-character strings terminated by a <code>wchar_t NULL</code>.</p>
wcscat(), wscat()	<p>The <code>wcscat()</code> and <code>wscat()</code> functions append a copy of the wide-character string pointed to by <i>ws2</i> (including the terminating null wide-character code) to the end of the wide-character string pointed to by <i>ws1</i>. The initial wide-character code of <i>ws2</i> overwrites the null wide-character code at the end of <i>ws1</i>. If copying takes place between objects that overlap, the behavior is undefined. Both functions return <i>s1</i>; no return value is reserved to indicate an error.</p>
wcsncat(), wsnocat()	<p>The <code>wcsncat()</code> and <code>wsncat()</code> functions append not more than <i>n</i> wide-character codes (a null wide-character code and wide-character codes that follow it are not appended) from the array pointed to by <i>ws2</i> to the end of the wide-character string pointed to by <i>ws1</i>. The initial wide-character code of <i>ws2</i> overwrites the null wide-character code at the end of <i>ws1</i>. A terminating null wide-character code is always appended to the result. Both functions return <i>ws1</i>; no return value is reserved to indicate an error.</p>
wcscmp(), wscmp()	<p>The <code>wcscmp()</code> and <code>wscmp()</code> functions compare the wide-character string pointed to by <i>ws1</i> to the wide-character string pointed to by <i>ws2</i>. The sign of a non-zero return value is determined by the sign of the difference between the values of the first pair of wide-character codes that differ in the objects being compared. Upon completion, both functions return an integer greater than, equal to, or less than zero, if the wide-character string pointed to by <i>ws1</i> is greater than, equal to, or less than the wide-character string pointed to by <i>ws2</i>.</p>

wcstring(3C)

wcsncmp() , wsncmp()	The <code>wcsncmp()</code> and <code>wsncmp()</code> functions compare not more than <i>n</i> wide-character codes (wide-character codes that follow a null wide character code are not compared) from the array pointed to by <i>ws1</i> to the array pointed to by <i>ws2</i> . The sign of a non-zero return value is determined by the sign of the difference between the values of the first pair of wide-character codes that differ in the objects being compared. Upon successful completion, both functions return an integer greater than, equal to, or less than zero, if the possibly null-terminated array pointed to by <i>ws1</i> is greater than, equal to, or less than the possibly null-terminated array pointed to by <i>ws2</i> .
wscpy() , wscopy()	The <code>wscopy()</code> and <code>wscpy()</code> functions copy the wide-character string pointed to by <i>ws2</i> (including the terminating null wide-character code) into the array pointed to by <i>ws1</i> . If copying takes place between objects that overlap, the behavior is undefined. Both functions return <i>ws1</i> ; no return value is reserved to indicate an error.
wcsncpy() , wsncpy()	The <code>wcsncpy()</code> and <code>wsncpy()</code> functions copy not more than <i>n</i> wide-character codes (wide-character codes that follow a null wide character code are not copied) from the array pointed to by <i>ws2</i> to the array pointed to by <i>ws1</i> . If copying takes place between objects that overlap, the behavior is undefined. If the array pointed to by <i>ws2</i> is a wide-character string that is shorter than <i>n</i> wide-character codes, null wide-character codes are appended to the copy in the array pointed to by <i>ws1</i> , until a total <i>n</i> wide-character codes are written. Both functions return <i>ws1</i> ; no return value is reserved to indicate an error.
wcslen() , wslen()	The <code>wcslen()</code> and <code>wslen()</code> functions compute the number of wide-character codes in the wide-character string to which <i>ws</i> points, not including the terminating null wide-character code. Both functions return <i>ws</i> ; no return value is reserved to indicate an error.
wcschr() , wschr()	The <code>wcschr()</code> and <code>wschr()</code> functions locate the first occurrence of <i>wc</i> in the wide-character string pointed to by <i>ws</i> . The value of <i>wc</i> must be a character representable as a type <code>wchar_t</code> and must be a wide-character code corresponding to a valid character in the current locale. The terminating null wide-character code is considered to be part of the wide-character string. Upon completion, both functions return a pointer to the wide-character code, or a null pointer if the wide-character code is not found.
wcsrchr() , wsrchr()	The <code>wcsrchr()</code> and <code>wsrchr()</code> functions locate the last occurrence of <i>wc</i> in the wide-character string pointed to by <i>ws</i> . The value of <i>wc</i> must be a character representable as a type <code>wchar_t</code> and must be a wide-character code corresponding to a valid character in the current locale. The terminating null wide-character code is considered to be part of the wide-character string. Upon successful completion, both functions return a pointer to the wide-character code, or a null pointer if <i>wc</i> does not occur in the wide-character string.
windex() , wrindex()	The <code>windex()</code> and <code>wrindex()</code> functions behave the same as <code>wschr()</code> and <code>wsrchr()</code> , respectively.

wcspbrk(), wspbrk()	The <code>wcspbrk()</code> and <code>wspbrk()</code> functions locate the first occurrence in the wide character string pointed to by <i>ws1</i> of any wide-character code from the wide-character string pointed to by <i>ws2</i> . Upon successful completion, the function returns a pointer to the wide-character code, or a null pointer if no wide-character code from <i>ws2</i> occurs in <i>ws1</i> .
wcswcs()	The <code>wcswcs()</code> function locates the first occurrence in the wide-character string pointed to by <i>ws1</i> of the sequence of wide-character codes (excluding the terminating null wide-character code) in the wide-character string pointed to by <i>ws2</i> . Upon successful completion, the function returns a pointer to the located wide-character string, or a null pointer if the wide-character string is not found. If <i>ws2</i> points to a wide-character string with zero length, the function returns <i>ws1</i> .
wcsspn(), wsspnl()	The <code>wcsspn()</code> and <code>wsspnl()</code> functions compute the length of the maximum initial segment of the wide-character string pointed to by <i>ws1</i> which consists entirely of wide-character codes from the wide-character string pointed to by <i>ws2</i> . Both functions return the length <i>ws1</i> ; no return value is reserved to indicate an error.
wcscspnl(), wscspnl()	The <code>wcscspnl()</code> and <code>wscspnl()</code> functions compute the length of the maximum initial segment of the wide-character string pointed to by <i>ws1</i> which consists entirely of wide-character codes <i>not</i> from the wide-character string pointed to by <i>ws2</i> . Both functions return the length of the initial substring of <i>ws1</i> ; no return value is reserved to indicate an error.
wcstok(), wstok()	A sequence of calls to the <code>wcstok()</code> and <code>wstok()</code> functions break the wide-character string pointed to by <i>ws1</i> into a sequence of tokens, each of which is delimited by a wide-character code from the wide-character string pointed to by <i>ws2</i> .
Default and other standards	<p>The third argument points to a caller-provided <code>wchar_t</code> pointer into which the <code>wcstok()</code> function stores information necessary for it to continue scanning the same wide-character string. This argument is not available with the XPG4 and SUS versions of <code>wcstok()</code>, nor is it available with the <code>wstok()</code> function. See <code>standards(5)</code>.</p> <p>The first call in the sequence has <i>ws1</i> as its first argument, and is followed by calls with a null pointer as their first argument. The separator string pointed to by <i>ws2</i> may be different from call to call.</p> <p>The first call in the sequence searches the wide-character string pointed to by <i>ws1</i> for the first wide-character code that is <i>not</i> contained in the current separator string pointed to by <i>ws2</i>. If no such wide-character code is found, then there are no tokens in the wide-character string pointed to by <i>ws1</i>, and <code>wcstok()</code> and <code>wstok()</code> return a null pointer. If such a wide-character code is found, it is the start of the first token.</p> <p>The <code>wcstok()</code> and <code>wstok()</code> functions then search from that point for a wide-character code that <i>is</i> contained in the current separator string. If no such wide-character code is found, the current token extends to the end of the wide-character string pointed to by <i>ws1</i>, and subsequent searches for a token will return a null pointer. If such a wide-character code is found, it is overwritten by a null wide character, which terminates the current token. The <code>wcstok()</code> and <code>wstok()</code> functions save a pointer to the following wide-character code, from which the next search for a token will start.</p>

wcstring(3C)

Each subsequent call, with a null pointer as the value of the first argument, starts searching from the saved pointer and behaves as described above.

Upon successful completion, both functions return a pointer to the first wide-character code of a token. Otherwise, if there is no token, a null pointer is returned.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	See NOTES.
MT-Level	MT-Safe

SEE ALSO `malloc(3C)`, `string(3C)`, `wcswidth(3C)`, `wcwidth(3C)`, `attributes(5)`, `standards(5)`

NOTES The `wscat()`, `wsnecat()`, `wscmp()`, `wsncmp()`, `wscopy()`, `wsncpy()`, `wcslen()`, `wchr()`, `wsrchr()`, `wspbrk()`, `wcswcs()`, `wcsspn()`, `wcscspn()`, and `wstok()` functions are Standard. The `wscat()`, `wsnecat()`, `wscmp()`, `wsncmp()`, `wscopy()`, `wsncpy()`, `wcslen()`, `wchr()`, `wsrchr()`, `wspbrk()`, `wcsspn()`, `wstok()`, `windex()`, and `wrindex()` functions are Stable.

NAME	wcswidth – number of column positions of a wide-character string
SYNOPSIS	<pre>#include <wchar.h> int wcswidth(const wchar_t *pwcs, size_t n);</pre>
DESCRIPTION	The <code>wcswidth()</code> function determines the number of column positions required for <i>n</i> wide-character codes (or fewer than <i>n</i> wide-character codes if a null wide-character code is encountered before <i>n</i> wide-character codes are exhausted) in the string pointed to by <i>pwcs</i> .
RETURN VALUES	The <code>wcswidth()</code> function either returns 0 (if <i>pwcs</i> points to a null wide-character code), or returns the number of column positions to be occupied by the wide-character string pointed to by <i>pwcs</i> , or returns -1 (if any of the first <i>n</i> wide-character codes in the wide-character string pointed to by <i>pwcs</i> is not a printing wide-character code).
ERRORS	No errors are defined.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

SEE ALSO [setlocale\(3C\)](#), [wctype\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

wcsxfrm(3C)

NAME	wcsxfrm, wsxfrm – wide character string transformation				
SYNOPSIS	<pre>#include <wchar.h> size_t wcsxfrm(wchar_t *ws1, const wchar_t *ws2, size_t n); size_t wsxfrm(wchar_t *ws1, const wchar_t *ws2, size_t n);</pre>				
DESCRIPTION	<p>The <code>wcsxfrm()</code> and <code>wsxfrm()</code> functions transform the wide character string pointed to by <code>ws2</code> and place the resulting wide character string into the array pointed to by <code>ws1</code>. The transformation is such that if either the <code>wcscmp(3C)</code> or <code>wscmp(3C)</code> functions are applied to two transformed wide strings, they return a value greater than, equal to, or less than 0, corresponding to the result of the <code>wscoll(3C)</code> or <code>wscoll(3C)</code> function applied to the same two original wide character strings. No more than <code>n</code> wide-character codes are placed into the resulting array pointed to by <code>ws1</code>, including the terminating null wide-character code. If <code>n</code> is 0, <code>ws1</code> is permitted to be a null pointer. If copying takes place between objects that overlap, the behavior is undefined.</p> <p>The <code>wcsxfrm()</code> and <code>wsxfrm()</code> functions do not change the setting of <code>errno</code> if successful.</p> <p>Since no return value is reserved to indicate an error, an application wishing to check for error situations should set <code>errno</code> to 0, then call <code>wcsxfrm()</code> or <code>wsxfrm()</code>, then check <code>errno</code>.</p>				
RETURN VALUES	<p>The <code>wcsxfrm()</code> and <code>wsxfrm()</code> functions return the length of the transformed wide character string (not including the terminating null wide-character code). If the value returned is <code>n</code> or more, the contents of the array pointed to by <code>ws1</code> are indeterminate.</p> <p>On error, <code>wcsxfrm()</code> and <code>wsxfrm()</code> may set <code>errno</code> but no return value is reserved to indicate an error.</p>				
ERRORS	<p>The <code>wcsxfrm()</code> and <code>wsxfrm()</code> functions may fail if:</p> <p>EINVAL The wide character string pointed to by <code>ws2</code> contains wide-character codes outside the domain of the collating sequence.</p>				
USAGE	<p>The transformation function is such that two transformed wide character strings can be ordered by the <code>wcscmp()</code> or <code>wscmp()</code> functions as appropriate to collating sequence information in the program's locale (category <code>LC_COLLATE</code>).</p> <p>The fact that when <code>n</code> is 0, <code>ws1</code> is permitted to be a null pointer, is useful to determine the size of the <code>ws1</code> array prior to making the transformation.</p>				
ATTRIBUTES	<p>See <code>attributes(5)</code> 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>CSI</td> <td>Enabled</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
CSI	Enabled				

Interface Stability	wcsxfrm() is Standard
MT-Level	MT-Safe with exceptions

The `wcsxfrm()` and `wsxfrm()` functions can be used safely in multithreaded applications as long as `setlocale(3C)` is not being called to change the locale.

SEE ALSO

`setlocale(3C)`, `wscmp(3C)`, `wscoll(3C)`, `wscmp(3C)`, `wscoll(3C)`, `attributes(5)`, `standards(5)`

wctob(3C)

NAME | wctob – wide-character to single-byte conversion

SYNOPSIS | #include <stdio.h>
#include <wchar.h>

int **wctob**(wint_t *c*);

DESCRIPTION | The `wctob()` function determines whether *c* corresponds to a member of the extended character set whose character representation is a single byte when in the initial shift state.

The behavior of this function is affected by the `LC_CTYPE` category of the current locale. See `environ(5)`.

RETURN VALUES | The `wctob()` function returns EOF if *c* does not correspond to a character with length one in the initial shift state. Otherwise, it returns the single-byte representation of that character.

ERRORS | No errors are defined.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

SEE ALSO | `btowc(3C)`, `setlocale(3C)`, `attributes(5)`, `environ(5)`, `standards(5)`

NOTES | The `wctob()` function can be used safely in multithreaded applications, as long as `setlocale(3C)` is not being called to change the locale.

NAME	wctomb – convert a wide-character code to a character								
SYNOPSIS	<pre>#include <stdlib.h> int wctomb(char *s, wchar_t <i>wchar</i>);</pre>								
DESCRIPTION	<p>The <code>wctomb()</code> function determines the number of bytes needed to represent the character corresponding to the wide-character code whose value is <code>wchar</code>. It stores the character representation (possibly multiple bytes) in the array object pointed to by <code>s</code> (if <code>s</code> is not a null pointer). At most <code>MB_CUR_MAX</code> bytes are stored.</p> <p>A call with <code>s</code> as a null pointer causes this function to return 0. The behavior of this function is affected by the <code>LC_CTYPE</code> category of the current locale.</p>								
RETURN VALUES	<p>If <code>s</code> is a null pointer, <code>wctomb()</code> returns 0 value. If <code>s</code> is not a null pointer, <code>wctomb()</code> returns <code>-1</code> if the value of <code>wchar</code> does not correspond to a valid character, or returns the number of bytes that constitute the character corresponding to the value of <code>wchar</code>.</p> <p>In no case will the value returned be greater than the value of the <code>MB_CUR_MAX</code> macro.</p>								
ERRORS	No errors are defined.								
USAGE	The <code>wctomb()</code> function can be used safely in a multithreaded application, as long as <code>setlocale(3C)</code> is not being called to change the locale.								
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:								
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>CSI</td> <td>Enabled</td> </tr> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	Standard	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
CSI	Enabled								
Interface Stability	Standard								
MT-Level	MT-Safe with exceptions								
SEE ALSO	<code>mblen(3C)</code> , <code>mbstowcs(3C)</code> , <code>mbtowc(3C)</code> , <code>setlocale(3C)</code> , <code>wcstombs(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>								

wctrans(3C)

NAME	wctrans – define character mapping								
SYNOPSIS	<pre>#include <wctype.h> wctrans_t wctrans(const char *charclass);</pre>								
DESCRIPTION	<p>The <code>wctrans()</code> function is defined for valid character mapping names identified in the current locale. The <i>charclass</i> is a string identifying a generic character mapping name for which codeset-specific information is required. The following character mapping names are defined in all locales – "tolower" and "toupper".</p> <p>The function returns a value of type <code>wctrans_t</code>, which can be used as the second argument to subsequent calls of <code>towctrans(3C)</code>. The <code>wctrans()</code> function determines values of <code>wctrans_t</code> according to the rules of the coded character set defined by character mapping information in the program's locale (category <code>LC_CTYPE</code>). The values returned by <code>wctrans()</code> are valid until a call to <code>setlocale(3C)</code> that modifies the category <code>LC_CTYPE</code>.</p>								
RETURN VALUES	The <code>wctrans()</code> function returns 0 if the given character mapping name is not valid for the current locale (category <code>LC_CTYPE</code>), otherwise it returns a non-zero object of type <code>wctrans_t</code> that can be used in calls to <code>towctrans(3C)</code> .								
ERRORS	The <code>wctrans()</code> function may fail if: EINVAL The character mapping name pointed to by <i>charclass</i> is not valid in the current locale.								
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:								
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>CSI</td><td>Enabled</td></tr><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe with exceptions</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	Standard	MT-Level	MT-Safe with exceptions
ATTRIBUTE TYPE	ATTRIBUTE VALUE								
CSI	Enabled								
Interface Stability	Standard								
MT-Level	MT-Safe with exceptions								
SEE ALSO	<code>setlocale(3C)</code> , <code>towctrans(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>								

NAME	wctype – define character class												
SYNOPSIS	<pre>#include <wchar.h> wctype_t wctype(const char *charclass) ;</pre>												
DESCRIPTION	<p>The <code>wctype()</code> function is defined for valid character class names as defined in the current locale. The <i>charclass</i> is a string identifying a generic character class for which codeset-specific type information is required. The following character class names are defined in all locales:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td>alnum</td> <td>alpha</td> <td>blank</td> </tr> <tr> <td>cntrl</td> <td>digit</td> <td>graph</td> </tr> <tr> <td>lower</td> <td>print</td> <td>punct</td> </tr> <tr> <td>space</td> <td>upper</td> <td>xdigit</td> </tr> </table> <p>Additional character class names defined in the locale definition file (category <code>LC_CTYPE</code>) can also be specified.</p> <p>The function returns a value of type <code>wctype_t</code>, which can be used as the second argument to subsequent calls of <code>iswctype(3C)</code>. <code>wctype()</code> determines values of <code>wctype_t</code> according to the rules of the coded character set defined by character type information in the program's locale (category <code>LC_CTYPE</code>). The values returned by <code>wctype()</code> are valid until a call to <code>setlocale(3C)</code> that modifies the category <code>LC_CTYPE</code>.</p>	alnum	alpha	blank	cntrl	digit	graph	lower	print	punct	space	upper	xdigit
alnum	alpha	blank											
cntrl	digit	graph											
lower	print	punct											
space	upper	xdigit											
RETURN VALUES	The <code>wctype()</code> function returns 0 if the given character class name is not valid for the current locale (category <code>LC_CTYPE</code>); otherwise it returns an object of type <code>wctype_t</code> that can be used in calls to <code>iswctype()</code> .												
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:												
	<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>CSI</td> <td>Enabled</td> </tr> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe with exceptions</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	CSI	Enabled	Interface Stability	Standard	MT-Level	MT-Safe with exceptions				
ATTRIBUTE TYPE	ATTRIBUTE VALUE												
CSI	Enabled												
Interface Stability	Standard												
MT-Level	MT-Safe with exceptions												
SEE ALSO	<code>iswctype(3C)</code> , <code>setlocale(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>												

wcwidth(3C)

NAME	wcwidth – number of column positions of a wide-character code
SYNOPSIS	<pre>#include <wchar.h> int wcwidth(wchar_t <i>wc</i>);</pre>
DESCRIPTION	The <code>wcwidth()</code> function determines the number of column positions required for the wide character <code>wc</code> . The value of <code>wc</code> must be a character representable as a <code>wchar_t</code> , and must be a wide-character code corresponding to a valid character in the current locale.
RETURN VALUES	The <code>wcwidth()</code> function either returns 0 (if <code>wc</code> is a null wide-character code), or returns the number of column positions to be occupied by the wide-character code <code>wc</code> , or returns -1 (if <code>wc</code> does not correspond to a printing wide-character code).
ERRORS	No errors are defined.
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
CSI	Enabled
Interface Stability	Standard
MT-Level	MT-Safe with exceptions

SEE ALSO [setlocale\(3C\)](#), [wcswidth\(3C\)](#), [attributes\(5\)](#), [standards\(5\)](#)

NAME	wmemchr – find a wide-character in memory						
SYNOPSIS	<pre>#include <wchar.h> wchar_t *wmemchr(const wchar_t *ws, wchar_t wc, size_t n);</pre>						
ISO C++	<pre>#include <wchar.h> const wchar_t *wmemchr(const wchar_t *ws, wchar_t wc, size_t n); #include <cwchar> wchar_t *std::wmemchr(wchar_t *ws, wchar_t wc, size_t n);</pre>						
DESCRIPTION	<p>The <code>wmemchr()</code> function locates the first occurrence of <code>wc</code> in the initial <code>n</code> wide-characters of the object pointed to be <code>ws</code>. This function is not affected by locale and all <code>wchar_t</code> values are treated identically. The null wide-character and <code>wchar_t</code> values not corresponding to valid characters are not treated specially.</p> <p>If <code>n</code> is 0, <code>ws</code> must be a valid pointer and the function behaves as if no valid occurrence of <code>wc</code> is found.</p>						
RETURN VALUES	The <code>wmemchr()</code> function returns a pointer to the located wide-character, or a null pointer if the wide-character does not occur in the object.						
ERRORS	No errors are defined.						
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>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	wmemcmp(3C) , wmemcpy(3C) , wmemmove(3C) , wmemset(3C) , attributes(5) , standards(5)						

wmemcmp(3C)

NAME	wmemcmp – compare wide-characters in memory						
SYNOPSIS	<pre>#include <wchar.h> int wmemcmp(const wchar_t *ws1, const wchar_t *ws2, size_t n);</pre>						
DESCRIPTION	<p>The <code>wmemcmp()</code> function compares the first <i>n</i> wide-characters of the object pointed to by <i>ws1</i> to the first <i>n</i> wide-characters of the object pointed to by <i>ws2</i>. This function is not affected by locale and all <code>wchar_t</code> values are treated identically. The null wide-character and <code>wchar_t</code> values not corresponding to valid characters are not treated specially.</p> <p>If <i>n</i> is zero, <i>ws1</i> and <i>ws2</i> must be a valid pointers and the function behaves as if the two objects compare equal.</p>						
RETURN VALUES	The <code>wmemcmp()</code> function returns an integer greater than, equal to, or less than 0, accordingly as the object pointed to by <i>ws1</i> is greater than, equal to, or less than the object pointed to by <i>ws2</i> .						
ERRORS	No errors are defined.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>wmemchr(3C)</code> , <code>wmemcpy(3C)</code> , <code>wmemmove(3C)</code> , <code>wmemset(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	wmemcpy – copy wide-characters in memory						
SYNOPSIS	<pre>#include <wchar.h> wchar_t *wmemcpy(wchar_t *ws1, const wchar_t *ws2, size_t n);</pre>						
DESCRIPTION	<p>The <code>wmemcpy()</code> function copies <i>n</i> wide-characters from the object pointed to by <i>ws2</i> to the object pointed to be <i>ws1</i>. This function is not affected by locale and all <code>wchar_t</code> values are treated identically. The null wide-character and <code>wchar_t</code> values not corresponding to valid characters are not treated specially.</p> <p>If <i>n</i> is zero, <i>ws1</i> and <i>ws2</i> must be a valid pointers, and the function copies zero wide-characters.</p>						
RETURN VALUES	The <code>wmemcpy()</code> function returns the value of <i>ws1</i> .						
ERRORS	No errors are defined.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>wmemchr(3C)</code> , <code>wmemcmp(3C)</code> , <code>wmemmove(3C)</code> , <code>wmemset(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

wmemmove(3C)

NAME	wmemmove – copy wide-characters in memory with overlapping areas						
SYNOPSIS	<pre>#include <wchar.h> wchar_t *wmemmove(wchar_t *ws1, const wchar_t *ws2, size_t n);</pre>						
DESCRIPTION	<p>The <code>wmemmove()</code> function copies n wide-characters from the object pointed to by <code>ws2</code> to the object pointed to by <code>ws1</code>. Copying takes place as if the n wide-characters from the object pointed to by <code>ws2</code> are first copied into a temporary array of n wide-characters that does not overlap the objects pointed to by <code>ws1</code> or <code>ws2</code>, and then the n wide-characters from the temporary array are copied into the object pointed to by <code>ws1</code>.</p> <p>This function is not affected by locale and all <code>wchar_t</code> values are treated identically. The null wide-character and <code>wchar_t</code> values not corresponding to valid characters are not treated specially.</p> <p>If n is 0, <code>ws1</code> and <code>ws2</code> must be a valid pointers, and the function copies zero wide-characters.</p>						
RETURN VALUES	The <code>wmemmove()</code> function returns the value of <code>ws1</code> .						
ERRORS	No errors are defined.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Interface Stability</td><td>Standard</td></tr><tr><td>MT-Level</td><td>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>wmemchr(3C)</code> , <code>wmemcmp(3C)</code> , <code>wmemcpy(3C)</code> , <code>wmemset(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

NAME	wmemset – set wide-characters in memory						
SYNOPSIS	<pre>#include <wchar.h> wchar_t *wmemset(wchar_t *ws, wchar_t wc, size_t n);</pre>						
DESCRIPTION	<p>The <code>wmemset()</code> function copies the value of <code>wc</code> into each of the first <code>n</code> wide-characters of the object pointed to by <code>ws</code>. This function is not affected by locale and all <code>wchar_t</code> values are treated identically. The null wide-character and <code>wchar_t</code> values not corresponding to valid characters are not treated specially.</p> <p>If <code>n</code> is 0, <code>ws</code> must be a valid pointer and the function copies zero wide-characters.</p>						
RETURN VALUES	The <code>wmemset()</code> functions returns the value of <code>ws</code> .						
ERRORS	No errors are defined.						
ATTRIBUTES	See <code>attributes(5)</code> for descriptions of the following attributes:						
	<table border="1"> <thead> <tr> <th>ATTRIBUTE TYPE</th> <th>ATTRIBUTE VALUE</th> </tr> </thead> <tbody> <tr> <td>Interface Stability</td> <td>Standard</td> </tr> <tr> <td>MT-Level</td> <td>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Interface Stability	Standard	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Interface Stability	Standard						
MT-Level	MT-Safe						
SEE ALSO	<code>wmemchr(3C)</code> , <code>wmemcmp(3C)</code> , <code>wmemcpy(3C)</code> , <code>wmemmove(3C)</code> , <code>attributes(5)</code> , <code>standards(5)</code>						

wordexp(3C)

NAME	wordexp, wordfree – perform word expansions						
SYNOPSIS	<pre>#include <wordexp.h> int wordexp(const char *restrict <i>words</i>, wordexp_t *restrict <i>pwordexp</i>, int <i>flags</i>); void wordfree(wordexp_t *<i>pwordexp</i>);</pre>						
DESCRIPTION	<p>The <code>wordexp()</code> function performs word expansions, subject to quoting, and places the list of expanded words into the structure pointed to by <i>pwordexp</i>.</p> <p>The <code>wordfree()</code> function frees any memory allocated by <code>wordexp()</code> associated with <i>pwordexp</i>.</p>						
<i>words</i> Argument	<p>The <i>words</i> argument is a pointer to a string containing one or more words to be expanded. The expansions will be the same as would be performed by the shell if <i>words</i> were the part of a command line representing the arguments to a utility. Therefore, <i>words</i> must not contain an unquoted NEWLINE or any of the unquoted shell special characters:</p> <pre> & ; < ></pre> <p>except in the context of command substitution. It also must not contain unquoted parentheses or braces, except in the context of command or variable substitution. If the argument <i>words</i> contains an unquoted comment character (number sign) that is the beginning of a token, <code>wordexp()</code> may treat the comment character as a regular character, or may interpret it as a comment indicator and ignore the remainder of <i>words</i>.</p>						
<i>pwordexp</i> Argument	<p>The structure type <code>wordexp_t</code> is defined in the header <code><wordexp.h></code> and includes at least the following members:</p> <table><tr><td><code>size_t we_wordc</code></td><td>Count of words matched by <i>words</i>.</td></tr><tr><td><code>char **we_wordv</code></td><td>Pointer to list of expanded words.</td></tr><tr><td><code>size_t we_offs</code></td><td>Slots to reserve at the beginning of <i>pwordexp->we_wordv</i>.</td></tr></table> <p>The <code>wordexp()</code> function stores the number of generated words into <i>pwordexp->we_wordc</i> and a pointer to a list of pointers to words in <i>pwordexp->we_wordv</i>. Each individual field created during field splitting is a separate word in the <i>pwordexp->we_wordv</i> list. The words are in order. The first pointer after the last word pointer will be a null pointer.</p> <p>It is the caller's responsibility to allocate the storage pointed to by <i>pwordexp</i>. The <code>wordexp()</code> function allocates other space as needed, including memory pointed to by <i>pwordexp->we_wordv</i>. The <code>wordfree()</code> function frees any memory associated with <i>pwordexp</i> from a previous call to <code>wordexp()</code>.</p>	<code>size_t we_wordc</code>	Count of words matched by <i>words</i> .	<code>char **we_wordv</code>	Pointer to list of expanded words.	<code>size_t we_offs</code>	Slots to reserve at the beginning of <i>pwordexp->we_wordv</i> .
<code>size_t we_wordc</code>	Count of words matched by <i>words</i> .						
<code>char **we_wordv</code>	Pointer to list of expanded words.						
<code>size_t we_offs</code>	Slots to reserve at the beginning of <i>pwordexp->we_wordv</i> .						

flags Argument

The *flags* argument is used to control the behavior of `wordexp()`. The value of *flags* is the bitwise inclusive OR of zero or more of the following constants, which are defined in `<wordexp.h>`:

<code>WRDE_APPEND</code>	Append words generated to the ones from a previous call to <code>wordexp()</code> .
<code>WRDE_DOOFFS</code>	Make use of <code>pwordexp->we_offs</code> . If this flag is set, <code>pwordexp->we_offs</code> is used to specify how many NULL pointers to add to the beginning of <code>pwordexp->we_wordv</code> . In other words, <code>pwordexp->we_wordv</code> will point to <code>pwordexp->we_offs</code> NULL pointers, followed by <code>pwordexp->we_wordc</code> word pointers, followed by a NULL pointer.
<code>WRDE_NOCMD</code>	Fail if command substitution is requested.
<code>WRDE_REUSE</code>	The <code>pwordexp</code> argument was passed to a previous successful call to <code>wordexp()</code> , and has not been passed to <code>wordfree()</code> . The result will be the same as if the application had called <code>wordfree()</code> and then called <code>wordexp()</code> without <code>WRDE_REUSE</code> .
<code>WRDE_SHOWERR</code>	Do not redirect <code>stderr</code> to <code>/dev/null</code> .
<code>WRDE_UNDEF</code>	Report error on an attempt to expand an undefined shell variable.

The `WRDE_APPEND` flag can be used to append a new set of words to those generated by a previous call to `wordexp()`. The following rules apply when two or more calls to `wordexp()` are made with the same value of `pwordexp` and without intervening calls to `wordfree()`:

1. The first such call must not set `WRDE_APPEND`. All subsequent calls must set it.
2. All of the calls must set `WRDE_DOOFFS`, or all must not set it.
3. After the second and each subsequent call, `pwordexp->we_wordv` will point to a list containing the following:
 - a. zero or more NULL pointers, as specified by `WRDE_DOOFFS` and `pwordexp->we_offs`.
 - b. pointers to the words that were in the `pwordexp->we_wordv` list before the call, in the same order as before.
 - c. pointers to the new words generated by the latest call, in the specified order.
4. The count returned in `pwordexp->we_wordc` will be the total number of words from all of the calls.
5. The application can change any of the fields after a call to `wordexp()`, but if it does it must reset them to the original value before a subsequent call, using the same `pwordexp` value, to `wordfree()` or `wordexp()` with the `WRDE_APPEND` or `WRDE_REUSE` flag.

If *words* contains an unquoted:

wordexp(3C)

NEWLINE | & ; < > () { } in an inappropriate context, `wordexp()` will fail, and the number of expanded words will be zero.

Unless `WRDE_SHOWERR` is set in *flags*, `wordexp()` will redirect `stderr` to `/dev/null` for any utilities executed as a result of command substitution while expanding *words*.

If `WRDE_SHOWERR` is set, `wordexp()` may write messages to *stderr* if syntax errors are detected while expanding *words*. If `WRDE_DOOFFS` is set, then `pwordexp->we_offs` must have the same value for each `wordexp()` call and `wordfree()` call using a given *pwordexp*.

The following constants are defined as error return values:

<code>WRDE_BADCHAR</code>	One of the unquoted characters: NEWLINE & ; < > () { } appears in <i>words</i> in an inappropriate context.
<code>WRDE_BADVAL</code>	Reference to undefined shell variable when <code>WRDE_UNDEF</code> is set in <i>flags</i> .
<code>WRDE_CMDSUB</code>	Command substitution requested when <code>WRDE_NOCMD</code> was set in <i>flags</i> .
<code>WRDE_NOSPACE</code>	Attempt to allocate memory failed.
<code>WRDE_SYNTAX</code>	Shell syntax error, such as unbalanced parentheses or unterminated string.

RETURN VALUES On successful completion, `wordexp()` returns 0.

Otherwise, a non-zero value as described in `<wordexp.h>` is returned to indicate an error. If `wordexp()` returns the value `WRDE_NOSPACE`, then `pwordexp->we_wordc` and `pwordexp->we_wordv` will be updated to reflect any words that were successfully expanded. In other cases, they will not be modified.

The `wordfree()` function returns no value.

ERRORS No errors are defined.

USAGE This function is intended to be used by an application that wants to do all of the shell's expansions on a word or words obtained from a user. For example, if the application prompts for a filename (or list of filenames) and then uses `wordexp()` to process the input, the user could respond with anything that would be valid as input to the shell.

The `WRDE_NOCMD` flag is provided for applications that, for security or other reasons, want to prevent a user from executing shell command. Disallowing unquoted shell special characters also prevents unwanted side effects such as executing a command or writing a file.

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

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Standard
MT-Level	MT-Safe

SEE ALSO `fnmatch(3C)`, `glob(3C)`, `attributes(5)`, `standards(5)`

wsprintf(3C)

NAME	wsprintf – formatted output conversion				
SYNOPSIS	<pre>#include <stdio.h> #include <wchar.h> int wsprintf(wchar_t *s, const char *format, /* arg */ ...);</pre>				
DESCRIPTION	<p>The <code>wsprintf()</code> function outputs a Process Code string ending with a Process Code (<code>wchar_t</code>) null character. It is the user's responsibility to allocate enough space for this <code>wchar_t</code> string.</p> <p>This returns the number of Process Code characters (excluding the null terminator) that have been written. The conversion specifications and behavior of <code>wsprintf()</code> are the same as the regular sprintf(3C) function except that the result is a Process Code string for <code>wsprintf()</code>, and on Extended Unix Code (EUC) character string for <code>sprintf()</code>.</p>				
RETURN VALUES	Upon successful completion, <code>wsprintf()</code> returns the number of characters printed. Otherwise, a negative value is returned.				
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>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	wsscanf(3C) , printf(3C) , scanf(3C) , sprintf(3C) , attributes(5)				

NAME	wsscanf – formatted input conversion				
SYNOPSIS	<pre>#include<stdio.h> #include <wchar.h> int wsscanf(wchar_t *s, const char *format, /* pointer */ ...);</pre>				
DESCRIPTION	<p>The <code>wsscanf()</code> function reads Process Code characters from the Process Code string <code>s</code>, interprets them according to the <code>format</code>, and stores the results in its arguments. It expects, as arguments, a control string <code>format</code>, and a set of <code>pointer</code> arguments indicating where the converted input should be stored. The results are undefined if there are insufficient <code>args</code> for the format. If the format is exhausted while <code>args</code> remain, the excess <code>args</code> are simply ignored.</p> <p>The conversion specifications and behavior of <code>wsscanf()</code> are the same as the regular <code>sscanf(3C)</code> function except that the source is a Process Code string for <code>wsscanf()</code> and on Extended Unix Code (EUC) character string for <code>sscanf(3C)</code>.</p>				
RETURN VALUES	Upon successful completion, <code>wsscanf()</code> returns the number of characters matched. Otherwise, it returns a negative value.				
ATTRIBUTES	See <code>attributes(5)</code> 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>MT-Safe</td> </tr> </tbody> </table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	<code>wsprintf(3C)</code> , <code>printf(3C)</code> , <code>scanf(3C)</code> , <code>attributes(5)</code>				

wstring(3C)

NAME	wstring, wscasecmp, wsnccasecmp, wsdup, wscol – Process Code string operations				
SYNOPSIS	<pre>#include <wchar.h> int wscasecmp(const wchar_t *s1, const wchar_t *s2); int wsnccasecmp(const wchar_t *s1, const wchar_t *s2, int n); wchar_t *wsdup(const wchar_t *s); int wscol(const wchar_t *s);</pre>				
DESCRIPTION	<p>These functions operate on Process Code strings terminated by <code>wchar_t</code> null characters. During appending or copying, these routines do not check for an overflow condition of the receiving string. In the following, <code>s</code>, <code>s1</code>, and <code>s2</code> point to Process Code strings terminated by a <code>wchar_t</code> null.</p> <p>wscasecmp(), wsnccasecmp() The <code>wscasecmp()</code> function compares its arguments, ignoring case, and returns an integer greater than, equal to, or less than 0, depending upon whether <code>s1</code> is lexicographically greater than, equal to, or less than <code>s2</code>. It makes the same comparison but compares at most <code>n</code> Process Code characters. The four Extended Unix Code (EUC) codesets are ordered from lowest to highest as 0, 2, 3, 1 when characters from different codesets are compared.</p> <p>wsdup() The <code>wsdup()</code> function returns a pointer to a new Process Code string, which is a duplicate of the string pointed to by <code>s</code>. The space for the new string is obtained using malloc(3C). If the new string cannot be created, a null pointer is returned.</p> <p>wscol() The <code>wscol()</code> function returns the screen display width (in columns) of the Process Code string <code>s</code>.</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>MT-Safe</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	MT-Level	MT-Safe
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
MT-Level	MT-Safe				
SEE ALSO	malloc(3C) , string(3C) , wcstring(3C) , attributes(5)				

Index

A

- abort — terminate the process abnormally, 25
- abs — return absolute value of integer, 26
- access dynamic thread scheduling
 - thr_getprio, 828
 - thr_setprio, 828
- access dynamic thread scheduling parameters
 - pthread_getschedparam, 534
 - pthread_setschedparam, 534
- accounting, time accounting for current process
 - times, 845
- acquire and release stream lock — flockfile, 147
funlockfile, 147
- add close or open action to spawn file actions object —
 - posix_spawn_file_actions_addclose, 456
- add close or open action to spawn file actions object —
 - posix_spawn_file_actions_addopen, 456
- add dup2 action to spawn file actions object —
 - posix_spawn_file_actions_adddup2, 458
- add or change environment variable —
 - setenv, 671
- additional severities, define — addsev, 27
- address of symbol, get address in shared object or executable — dlsym, 109
- addsev — define additional severities, 27
- addseverity — build a list of severity levels for an application for use with fmtmsg, 28
- alarm, schedule signal after interval in microseconds — ualarm, 866
- allocate and deallocate process handles for libc_db — td_ta_delete, 799
- allocate and deallocate process handles for libc_db — td_ta_get_ph, 799
- allocate and deallocate process handles for libc_db — td_ta_new, 799
- allocation cache manipulation —
 - umem_cache_alloc, 875
- allocation cache manipulation —
 - umem_cache_create, 875
- allocation cache manipulation —
 - umem_cache_destroy, 875
- allocation cache manipulation —
 - umem_cache_free, 875
- allow signals to interrupt functions —
 - siginterrupt, 686
- alphasort — scan a directory, 644, 646
- applications
 - build a list of severity levels for use with fmtmsg — addseverity, 28
 - display a message on stderr or system console — fmtmsg, 149
 - get entries from symbol table — nlist, 412
- arithmetic, compute the quotient and remainder — div, 88
- arithmetic, 48-bit integer, generate uniformly distributed pseudo-random numbers —
 - drand48, 112
- ascftime — convert date and time to string, 722
- assert — verify program assertion, 30
- associate or dissociate the object with the port — port_associate, 427
- associate or dissociate the object with the port — port_dissociate, 427

associate a stream with a file descriptor —
 fdopen, 131

atexit — register a function to run at process
 termination or object unloading, 31

atof — convert string to floating-point
 number, 744

atoi — string conversion routines, 750

atol — string conversion routines, 750

atoll — string conversion routines, 750

atomic operations — atomic_add_16, 32

atomic operations — atomic_add_16_nv, 32

atomic operations — atomic_add_32, 32

atomic operations — atomic_add_32_nv, 32

atomic operations — atomic_add_64, 32

atomic operations — atomic_add_64_nv, 32

atomic operations — atomic_add_long, 32

atomic operations — atomic_add_long_nv, 32

atomic operations — atomic_and_32, 32

atomic operations — atomic_and_uint, 32

atomic operations — atomic_ops, 32

atomic operations — atomic_or_32, 32

atomic operations — atomic_or_uint, 32

atomic_add_16 — atomic operations, 32

atomic_add_16_nv — atomic operations, 32

atomic_add_32 — atomic operations, 32

atomic_add_32_nv — atomic operations, 32

atomic_add_64 — atomic operations, 32

atomic_add_64_nv — atomic operations, 32

atomic_add_long — atomic operations, 32

atomic_add_long_nv — atomic operations, 32

atomic_and_32 — atomic operations, 32

atomic_and_uint — atomic operations, 32

atomic_ops — atomic operations, 32

atomic_or_32 — atomic operations, 32

atomic_or_uint — atomic operations, 32

attach a STREAMS or doors-based file
 descriptor to an object in the file system
 name space — fattach, 123

atropen — open a file, 34

B

base-64 ASCII characters
 convert from long integer — l64a, 24
 convert to long integer — a64l, 24

basename — return the last element of path
 name, 35

bcmp — memory operations, 41

bcopy — memory operations, 41

binary input — fread, 169

binary output — fwrite, 199

binary search a sorted table — bsearch, 39

bind_textdomain_codeset — message handling
 functions, 272

bindtextdomain — message handling
 functions, 272

bit and byte operations, find first set bit —
 ffs, 136

bsd_signal — simplified signal facilities, 38

bsdmalloc — memory allocator, 36

bsearch — binary search a sorted table, 39

bstring — memory operations, 41

btowc — single-byte to wide-character
 conversion, 42

buffering, assign to stream
 — setbuffer, 669
 — setlinebuf, 669

byte swap — swab, 758

bzero — memory operations, 41

C

C Compilation

close a shared object — dlclose, 91

create new file from dynamic object
 component — dldump, 93

get address of symbol in shared object or
 executable — dlsym, 109

get diagnostic information — dlerror, 99

open a shared object — dlopen, 104

catclose — close a message catalog, 44

catgets — read a program message, 43

catopen — open a message catalog, 44

cfgetispeed — get input baud rate, 47

cfgetospeed — get output baud rate, 47

cfsetispeed — set input baud rate, 48

cfsetospeed — set output baud rate, 48

cftime — convert date and time to string, 722

change privilege sets and check whether
 privileges are set — priv_ineffect, 478

change privilege sets and check whether
 privileges are set — priv_set, 478

change or add value to environment —
 putenv, 600

change or examine calling thread's signal mask
 — pthread_sigmask, 586

change or examine calling thread's signal mask
 — thr_sigsetmask, 838

change the priority ceiling of a mutex
 — pthread_mutex_getprioceiling, 557
 — pthread_mutex_setprioceiling, 557

child processes
 allows a parent process to control the
 execution of a child process — ptrace, 596
 wait for child process to change state —
 waitpid, 908
 wait for child process to stop or terminate —
 wait, 902

clock — report CPU time used, 49

close or iterate over open file descriptors —
 closefrom, 51

close or iterate over open file descriptors —
 fdwalk, 51

close a directory stream — closedir, 50

close a shared object — dlclose, 91

close a stream — fclose, 127

closedir — close a directory stream, 50

closefrom — close or iterate over open file
 descriptors, 51

closelog — control system log, 769

code conversion allocation function —
 iconv_open, 313

code conversion deallocation function —
 iconv_close, 312

code conversion for Process Code and File Code
 — strtows, 755
 — wstostr, 755

code conversion function — iconv, 306

collect target process statistics for libc_db —
 td_ta_enable_stats, 790

collect target process statistics for libc_db —
 td_ta_get_stats, 790

collect target process statistics for libc_db —
 td_ta_reset_stats, 790

column positions of a wide-character code —
 wctype, 944

column positions of a wide-character string —
 wcsnwidth, 937

command option parsing — getopt, 233

command suboptions, parse suboptions from a
 string — getsubopt, 268

compare thread IDs — pthread_equal, 530

compare wide-characters in memory —
 wmemcmp, 946

compile and execute regular expressions
 — re_comp, 629
 — re_exec, 629

cond_broadcast — condition variables, 53

cond_destroy — condition variables, 53

cond_init — condition variables, 53
 Condition Signaling, 54
 Condition Wait, 54
 Destroy, 55
 Initialize, 53

cond_reltimedwait — condition variables, 53

cond_signal — condition variables, 53

cond_timedwait — condition variables, 53

cond_wait — condition variables, 53

condition variables — cond_broadcast, 53

condition variables — cond_destroy, 53

condition variables — cond_init, 53

condition variables — cond_reltimedwait, 53

condition variables — cond_signal, 53

condition variables — cond_timedwait, 53

condition variables — cond_wait, 53

confstr — get configurable variables, 58

control system log
 — closelog, 769
 — openlog, 769
 — setlogmask, 769
 — syslog, 769

convert a thread ID or LWP ID to a thread
 handle — td_ta_map_id2thr, 798

convert a thread ID or LWP ID to a thread
 handle — td_ta_map_lwp2thr, 798

convert date and time to string — strftime, 722

convert formatted input — fscanf, 647

convert formatted input — scanf, 647

convert formatted input — sscanf, 647

convert formatted input — vfscanf, 647

convert formatted input — vscanf, 647

convert formatted input — vsscanf, 647

convert formatted wide-character input —
 fwscanf, 200

convert formatted wide-character input —
 swscanf, 200

convert formatted wide-character input —
 vwscanf, 200

convert formatted wide-character input —
 vswscanf, 200

- convert formatted wide-character input — `vwscanf`, 200
- convert formatted wide-character input — `wscanf`, 200
- convert string to floating-point number — `atof`, 744
- convert string to floating-point number — `strtod`, 744
- convert string to floating-point number — `strtof`, 744
- convert string to floating-point number — `strtold`, 744
- convert string to unsigned long — `strtoul`, 753
- convert wide character string to floating-point number — `watof`, 922
- convert wide character string to floating-point number — `wcstod`, 922
- convert wide character string to floating-point number — `wcstof`, 922
- convert wide character string to floating-point number — `wcstold`, 922
- convert wide character string to floating-point number — `wstod`, 922
- convert wide character string to long integer — `watoi`, 926
- convert wide character string to long integer — `watol`, 926
- convert wide character string to long integer — `watoll`, 926
- convert wide character string to long integer — `wcstol`, 926
- convert wide character string to long integer — `wcstoll`, 926
- convert wide character string to long integer — `wstol`, 926
- convert wide character string to unsigned long — `wcstoul`, 930
- convert wide-character string to integer type — `wcstoimax`, 925
- convert wide-character string to integer type — `wcstoumax`, 925
- convert date and time to string — `strftime`
 - `ascftime`, 722
 - `cftime`, 722
- convert a character string to a wide-character string — `mbstowcs`, 371
- convert a character string to a wide-character string (restartable) — `mbsrtowcs`, 369
- convert a character to a wide-character code — `mbtowc`, 372
- convert a character to a wide-character code (restartable) — `mbrtowc`, 366
- convert a wide-character code to a character — `wctomb`, 941
- convert a wide-character code to a character (restartable) — `wcrtomb`, 915
- convert a wide-character string to a character string — `wcstombs`, 929
- convert a wide-character string to a character string (restartable) — `wcsrtoombs`, 919
- convert date and time to wide character string — `wcsftime`, 918
- convert floating-point number to string
 - `ecvt`, 117
 - `fcvt`, 117
 - `gcvt`, 117
- convert monetary value to string — `strfmon`, 718
- convert numbers to strings
 - `econvert`, 115
 - `fconvert`, 115
 - `fprintf`, 471
 - `gconvert`, 115
 - `printf`, 471
 - `qeconvert`, 115
 - `qfconvert`, 115
 - `qgconvert`, 115
 - `seconvert`, 115
 - `sfconvert`, 115
 - `sgconvert`, 115
 - `sprintf`, 471
 - `vfprintf`, 471
 - `vprintf`, 471
- convert to `wchar_t` strings, `wsprintf`, 954
- copy wide-characters in memory — `wmemcpy`, 947
- copy wide-characters in memory with overlapping areas — `wmemmove`, 948
- CPU time, report for calling process — `clock`, 49
- CPU-use, prepare execution profile — `monitor`, 388
- create a port — `port_create`, 431
- create a temporary file — `tmpfile`, 846
- create a thread — `thr_create`, 819
- create a thread — `pthread_create`, 525

create cancellation point in the calling thread. —
 pthread_testcancel, 595
 create new file from dynamic object component
 — dldump, 93
 create thread-specific data key —
 pthread_key_create, 539
 crypt — string encoding function, 64
 cset — get information on EUC codesets, 69
 csetcol — get information on EUC codesets, 69
 csetlen — get information on EUC codesets, 69
 csetno — get information on EUC codesets, 69
 ctermid — generate path name for controlling
 terminal, 70
 ctermid_r — generate path name for controlling
 terminal, 70
 current location of a named directory stream —
 telldir, 817
 current working directory, get pathname —
 getcwd, 208
 cuserid — get character-string representation of
 login name of user, 79

D

data base subroutines — dbm, 80
 dbmclose, 80
 dbmopen, 80
 delete, 80
 fetch, 80
 firstkey, 80
 nextkey, 80
 store, 80
 database functions — dbm_clearerr, 407
 database functions — dbm_close, 407
 database functions — dbm_delete, 407
 database functions — dbm_error, 407
 database functions — dbm_fetch, 407
 database functions — dbm_firstkey, 407
 database functions — dbm_nextkey, 407
 database functions — dbm_open, 407
 database functions — dbm_store, 407
 database functions — ndbm, 407
 date and time
 convert to string — asctime, 71
 convert to string — asctime_r, 71
 convert to string — ctime, 71
 convert to string — ctime_r, 71

date and time (Continued)
 convert to string — gmtime, 71
 convert to string — gmtime_r, 71
 convert to string — localtime, 71
 convert to string — localtime_r, 71
 convert to string — tzset, 71
 convert user format date and time —
 getdate, 210
 — gettimeofday, 276
 date and time conversion — strtptime, 737
 dbm — data base subroutines, 80
 dbm_clearerr — database functions, 407
 dbm_close — database functions, 407
 dbm_delete — database functions, 407
 dbm_error — database functions, 407
 dbm_fetch — database functions, 407
 dbm_firstkey — database functions, 407
 dbm_nextkey — database functions, 407
 dbm_open — database functions, 407
 dbm_store — database functions, 407
 dbmclose — data base subroutines, 80
 dbmopen — data base subroutines, 80
 dcgettext — message handling functions, 272
 dcngettext — message handling functions, 272
 debugging features of the umem library —
 umem_debug, 883
 debugging memory allocator
 — calloc, 912
 — cfree, 912
 — free, 912
 — mallinfo, 912
 — malloc, 912
 — mallopt, 912
 — memalign, 912
 — realloc, 912
 — valloc, 912
 decimal record from double-precision floating
 — double_to_decimal, 143
 decimal record from extended-precision floating
 — extended_to_decimal, 143
 decimal record from quadruple-precision
 floating — quadruple_to_decimal, 143
 decimal record from single-precision floating —
 single_to_decimal, 143
 decimal record to double-precision floating —
 decimal_to_double, 82
 decimal record to extended-precision floating —
 decimal_to_extended, 82

decimal record to quadruple-precision floating
 — `decimal_to_quadruple`, 82
 decimal record to single-precision floating —
 `decimal_to_single`, 82
`decimal_to_double` — decimal record to
 double-precision floating, 82
`decimal_to_extended` — decimal record to
 extended-precision floating, 82
`decimal_to_quadruple` — decimal record to
 quadruple-precision floating, 82
`decimal_to_single` — decimal record to
 single-precision floating, 82
 define character class — `wctype`, 943
 define character mapping — `wctrans`, 942
 define default catalog — `setcat`, 670
 define the label for `pfmt()` and `lfmt()`. —
 `setlabel`, 679
 delete — data base subroutines, 80
 delete thread-specific data key —
 `pthread_key_delete`, 541
 destroy and initialize a barrier object —
 `pthread_barrier_destroy`, 505
 destroy and initialize a barrier object —
 `pthread_barrier_init`, 505
 destroy and initialize barrier attributes object —
 `pthread_barrierattr_destroy`, 502
 destroy and initialize barrier attributes object —
 `pthread_barrierattr_init`, 502
 destroy and initialize spawn attributes object —
 `posix_spawnattr_destroy`, 447
 destroy and initialize spawn attributes object —
 `posix_spawnattr_init`, 447
 destroy and initialize spawn file actions object
 — `posix_spawn_file_actions_destroy`, 459
 destroy and initialize spawn file actions object
 — `posix_spawn_file_actions_init`, 459
 destroy or initialize a spin lock object —
 `pthread_spin_destroy`, 591
 destroy or initialize a spin lock object —
 `pthread_spin_init`, 591
 detach a name from a STREAMS-based file
 descriptor — `fdetach`, 129
 detach a thread — `pthread_detach`, 529
 determine type of floating-point number —
 `finite`, 323
 determine type of floating-point number —
 `fpclass`, 323
 determine type of floating-point number —
 `isnand`, 323
 determine type of floating-point number —
 `isnanf`, 323
 determine type of floating-point number —
 `unordered`, 323
 determine conversion object status —
 `mbsinit`, 368
 determine if address is within stack boundaries
 — `stack_inbounds`, 707
 determine stack boundary violation event —
 `stack_violation`, 709
 device number, manage — `makedev`, major,
 minor, 354
`dgettext` — message handling functions, 272
`difftime` — computes the difference between
 two calendar times, 83
`directio` — provide advice to file system, 84
 directories
 get current working directory pathname —
 `getwd`, 293
 get pathname of current working directory
 — `getcwd`, 208
 directory operations
 — `alphasort`, 646
 — `scandir`, 646
`dirname` — report parent directory name of file
 path name, 86
 display error message in standard format —
 `pfmt`, 418
 display error message in standard format and
 pass to logging and monitoring services —
 `lfmt`, 332
 display error message in standard format and
 pass to logging and monitoring services —
 `vlfmt`, 892
 display error message in standard format and
 pass to logging and monitoring services —
 `vpfmt`, 894
`div` — compute quotient and remainder, 88
 division and remainder operations
 — `div`, 88
 — `ldiv`, 88
`dladdr` — translate address to symbolic
 information, 89
`dladdr1` — translate address to symbolic
 information, 89
`dlclose` — close a shared object, 91

dldump — create new file from dynamic object component of calling process, 93
 dlerror — get diagnostic information, 99
 dlinfo — dynamic load information, 100
 dlopen — open a shared object, 104
 dlsym — get address of symbol in shared object or executable, 109
 dngettext — message handling functions, 272
 double_to_decimal — decimal record from double-precision floating, 143
 dup2 — duplicate an open file descriptor, 114
 duplicate an open file descriptor — dup2, 114
 dynamic linking
 close a shared object — dlclose, 91
 create new file from dynamic object component — dldump, 93
 get address of symbol in shared object or executable — dlsym, 109
 get diagnostic information — dlerror, 99
 open a shared object — dlopen, 104
 dynamic load information — dlinfo, 100
 dynamic thread scheduling parameters access — pthread_setschedprio, 585

E

econvert — convert number to ASCII, 115
 edata — last location in program, 120
 enable or disable cancellation — pthread_setcancelstate, 581
 enabling or disabling cancellation — pthread_setcancelstate, 581
 end — last location in program, 120
 endgrent — group database entry functions, 219
 endpwent — get password entry from user database, 255
 endspent — get shadow password database entry, 264
 endusershell — get legal user shells, 281
 endutent — user accounting database functions, 282
 endutxent — user accounting database functions, 285
 environment name, return value — getenv, 217
 error messages, system, print — perror, 417
 etext — last location in program, 120

EUC character bytes, — euclen, 121
 EUC characters, convert a string of Process Code characters to EUC characters and put it on a stream — putwts, 604
 EUC codeset, get information, — getwidth, 294
 EUC codesets, get information
 — cset, 69
 — csetcol, 69
 — csetlen, 69
 — csetno, 69
 — wcsetno, 69
 EUC display width
 — euccol, 121
 — eucscol, 121
 euccol — get EUC character display width, 121
 euclen — get EUC byte length, 121
 eucscol — get EUC string display width, 121
 exit — terminate process, 122
 express an intention to extend the stack — _stack_grow, 706
 extended_to_decimal — decimal record from extended-precision floating, 143
 Extended Unix Code, *See* EUC

F

fast, scalable memory allocation — umem_alloc, 870
 fast, scalable memory allocation — umem_free, 870
 fast, scalable memory allocation — umem_nofail_callback, 870
 fast, scalable memory allocation — umem_zalloc, 870
 fattach — attach a STREAMS or doors-based file descriptor to an object in the file system name space, 123
 __fbufsize — interfaces to stdio FILE structure, 125
 fclose — close a stream, 127
 fconvert — convert number to ASCII, 115
 FD_CLR — synchronous I/O multiplexing, 658
 FD_ISSET — synchronous I/O multiplexing, 658
 FD_SET — synchronous I/O multiplexing, 658
 FD_ZERO — synchronous I/O multiplexing, 658

fdetach — detach a name from a STREAMS-based file descriptor, 129
 fdopen — associate a stream with a file descriptor, 131
 fdopendir — open directory, 415
 fdwalk — close or iterate over open file descriptors, 51
 fetch — data base subroutines, 80
 fflush — flush a stream, 134
 ffs — find first set bit, 136
 fgetc — get a byte from a stream, 137
 fgetgrent — group database entry functions, 219
 fgetgrent_r — group database entry functions, 219
 fgetpos — get current file position information, 140
 fgetpwent — get password entry from a file, 255
 fgetpwent_r — get password entry from a file, 255
 fgetspent — get shadow password database entry, 264
 fgetspent_r — get shadow password database entry(reentrant), 264
 fgetwc — get a wide-character code from a stream, 141
 fgetws — get a wide-character string from a stream, 295
 file descriptor
 duplicate an open one — dup2, 114
 test for a STREAMS file — isastream, 321
 file descriptors, apply or remove advisory lock on open file — flock, 145
 file pointer in a stream, reposition — fsetpos, 177
 file_to_decimal — decimal record from character stream, 733
 file tree
 recursively descend — ftw, 184
 recursively descend — nftw, 184
 files
 allows sections of file to be locked — lockf, 342
 optimizing usage of files — directio, 84
 — remove, 638
 report parent directory of file path name — dirname, 86
 files (Continued)
 set a file to a specified length — truncate, 857
 find pathname of a terminal — ttyname, 863
 find pathname of a terminal — ttyname_r, 863
 find a wide-character in memory — wmemchr, 945
 find a wide-character substring — wcsstr, 921
 find pathname of a terminal
 — ttyname, 863
 — ttyname_r, 863
 finite — determine type of floating-point number, 323
 firstkey — data base subroutines, 80
 __flbf — interfaces to stdio FILE structure, 125
 floating-point number, convert to string — ecvt, 117
 flock — apply or remove an advisory lock on an open file, 145
 flockfile — acquire and release stream lock, 147
 flush a stream — fflush, 134
 flush non-transmitted output data, non-read input data or both — tflush, 776
 __flushlbf — interfaces to stdio FILE structure, 125
 fmtmsg — display a message on stderr or system console, 149
 fnmatch — match filename or path name, 154
 fopen — open a stream, 156
 fopen — open stream, 159
 formatted input conversion — wscanf, 955
 formatted output conversion
 — fprintf, 471
 — printf, 471
 — sprintf, 471
 — vfprintf, 471
 — vprintf, 471
 — vsprintf, 471
 fpclass — determine type of floating-point number, 323
 __fpending — interfaces to stdio FILE structure, 125
 fpgetmask — IEEE floating-point environment control, 161
 fpgetround — IEEE floating-point environment control, 161
 fpgetsticky — IEEE floating-point environment control, 161

fprintf — formatted output conversion, 471
 fprintf — print formatted output, 460
 fpsetmask — IEEE floating-point environment control, 161
 fpsetround — IEEE floating-point environment control, 161
 fpsetsticky — IEEE floating-point environment control, 161
 __fpurge — interfaces to stdio FILE structure, 125
 fputc — put a byte on a stream, 163
 fputwc — put wide-character code on a stream, 166
 fputws — put wide character string on a stream, 168
 fread — binary input, 169
 __freadable — interfaces to stdio FILE structure, 125
 __freading — interfaces to stdio FILE structure, 125
 free — memory allocator, 36
 freopen — open a stream, 171
 freopen — open stream, 159
 fscanf — convert formatted input, 647
 fseek — reposition a file-position indicator in a stream, 174
 fseeko — reposition a file-position indicator in a stream, 174
 __fsetlocking — interfaces to stdio FILE structure, 125
 fsetpos — reposition a file pointer in a stream, 177
 fsync — synchronize changes to a file, 178
 ftell — return a file offset in a stream, 180
 ftello — return a file offset in a stream, 180
 ftime — get date and time, 181
 ftruncate — set a file to a specified length, 857
 ftw — walk a file tree, 184
 func_to_decimal — decimal record from character function, 733
 funlockfile — acquire and release stream lock, 147
 fwide — set stream orientation, 189
 fwprintf — print formatted wide-character output, 190
 __fwritable — interfaces to stdio FILE structure, 125
 fwrite — binary output, 199

__fwriting — interfaces to stdio FILE structure, 125
 fwscanf — convert formatted wide-character input, 200

G

gconvert — convert number to ASCII, 115
 general terminal interface, — `termios`, 818
 generate path name for controlling terminal
 — `ctermid`, 70
 — `ctermid_r`, 70
 generate path names matching a pattern
 — `glob`, 298
 — `globfree`, 298
 get a wide-character string from a stream —
 `fgetws`, 295
 get a wide-character string from a stream —
 `getws`, 295
 get and set process-shared attribute of barrier
 attributes object —
 `pthread_barrierattr_getpshared`, 503
 get and set process-shared attribute of barrier
 attributes object —
 `pthread_barrierattr_setpshared`, 503
 get and set spawn-flags attribute of spawn
 attributes object —
 `posix_spawnattr_getflags`, 449
 get and set spawn-flags attribute of spawn
 attributes object —
 `posix_spawnattr_setflags`, 449
 get and set spawn-pgroup attribute of spawn
 attributes object —
 `posix_spawnattr_getpgroup`, 451
 get and set spawn-pgroup attribute of spawn
 attributes object —
 `posix_spawnattr_setpgroup`, 451
 get and set spawn-schedparam attribute of
 spawn attributes object —
 `posix_spawnattr_getschedparam`, 452
 get and set spawn-schedparam attribute of
 spawn attributes object —
 `posix_spawnattr_setschedparam`, 452
 get and set spawn-schedpolicy attribute of
 spawn attributes object —
 `posix_spawnattr_getschedpolicy`, 453

get and set spawn-schedpolicy attribute of spawn attributes object —
 posix_spawnattr_setschedpolicy, 453
 get and set spawn-sigdefault attribute of spawn attributes object —
 posix_spawnattr_getsigdefault, 454
 get and set spawn-sigdefault attribute of spawn attributes object —
 posix_spawnattr_setsigdefault, 454
 get and set spawn-sigmask attribute of spawn attributes object —
 posix_spawnattr_getsigmask, 455
 get and set spawn-sigmask attribute of spawn attributes object —
 posix_spawnattr_setsigmask, 455
 get and set stack attributes —
 pthread_attr_getstack, 496
 get and set stack attributes —
 pthread_attr_setstack, 496
 get and set the clock selection condition variable attribute — pthread_condattr_getclock, 512
 get and set the clock selection condition variable attribute — pthread_condattr_setclock, 512
 get error message string — strerror, 717
 get error message string — strerror_r, 717
 get legal user shells — endusershell, 281
 get legal user shells — getusershell, 281
 get legal user shells — setusershell, 281
 get mnttab file information —
 getextmntent, 229
 get mnttab file information — getmntany, 229
 get mnttab file information — getmntent, 229
 get mnttab file information — hasmntopt, 229
 get mnttab file information — putmntent, 229
 get mnttab file information — resetmnttab, 229
 get a byte from a stream
 — fgetc, 137
 — getc, 137
 — getc_unlocked, 137
 — getchar, 137
 — getchar_unlocked, 137
 — getw, 137
 get a synchronization object handle from a synchronization object's address —
 td_ta_map_addr2sync, 797
 get a thread's thread-specific data for libc_db library of interfaces — td_thr_tsd, 814
 get a wide-character code from a stream —
 fgetwc, 141
 get address of symbol in shared object or executable — dlsym, 109
 get an identifier for the current host —
 gethostid, 223
 get and set prioceiling attribute of mutex attribute object
 — pthread_mutexattr_getprioceiling, 543
 — pthread_mutexattr_setprioceiling, 543
 get and set process-shared attribute
 — pthread_mutexattr_getpshared, 548
 — pthread_mutexattr_setpshared, 548
 get and set process-shared attribute of read-write lock attributes object
 — pthread_rwlockattr_getpshared, 567
 — pthread_rwlockattr_setpshared, 567
 get and set protocol attribute of mutex attribute object
 — pthread_mutexattr_getprotocol, 545
 — pthread_mutexattr_setprotocol, 545
 get calling thread's ID — pthread_self, 580
 get calling thread's ID — thr_self, 837
 get configurable variables — confstr, 58
 get connected socket/STREAM peer's credentials — getpeerucred, 249
 get current file position information —
 fgetpos, 140
 get diagnostic information — dlerror, 99
 get foreground process group ID —
 tcgetpgrp, 778
 get input baud rate, — cfgetispeed, 47
 get name of signal — strsignal, 743
 get number of bytes in a character —
 mblen, 363
 get number of bytes in a character (restartable)
 — mbrlen, 364
 get or set a mutex type
 — pthread_mutexattr_gettype, 552
 — pthread_mutexattr_settype, 552
 get or set contentionscope attribute
 — pthread_attr_getscope, 495
 — pthread_attr_setscope, 495
 get or set detachstate attribute
 — pthread_attr_getdetachstate, 488
 — pthread_attr_setdetachstate, 488
 get or set inheritsched attribute
 — pthread_attr_getinheritsched, 491

get or set inheritsched attribute (Continued)
 — pthread_attr_setinheritsched, 491
 get or set level of concurrency
 — pthread_getconcurrency, 532
 — pthread_setconcurrency, 532
 get or set process scheduling priority
 — getpriority, 251
 — setpriority, 251
 get or set schedparam attribute
 — pthread_attr_getschedparam, 493
 — pthread_attr_setschedparam, 493
 get or set schedpolicy attribute
 — pthread_attr_getschedpolicy, 494
 — pthread_attr_setschedpolicy, 494
 get or set stackaddr attribute
 — pthread_attr_getstackaddr, 498
 — pthread_attr_setstackaddr, 498
 get or set stacksize attribute
 — pthread_attr_getstacksize, 499
 — pthread_attr_setstacksize, 499
 get or set the process-shared condition variable attributes
 — pthread_condattr_getpshared, 513
 — pthread_condattr_setpshared, 513
 get or set the thread guardsize attribute
 — pthread_attr_getguardsize, 489
 — pthread_attr_setguardsize, 489
 get output baud rate, — cfgetospeed, 47
 get process group ID for session leader for controlling terminal — tcgetsid, 779
 get system load averages — getloadavg, 226
 get system load averages for a processor set — pset_getloadavg, 483
 get the parameters associated with the terminal — tcgetattr, 777
 get thread information in libc_db library of interfaces — td_thr_get_info, 807
 get thread stack address and size — thr_stksegment, 842
 get wide character from a stream — getwc, 291
 get wide character from stdin stream — getwchar, 292
 getc — get a byte from a stream, 137
 getc_unlocked — get a byte from a stream, 137
 getchar — get a byte from a stream, 137
 getchar_unlocked — get a byte from a stream, 137
 getcpuid — obtain information on scheduling decisions, 207
 getcwd — get pathname of current working directory, 208
 getdate — convert user format date and time, 210
 General Specifications, 213
 Internal Format Conversion, 212
 Modified Conversion Specifications, 211
 getenv — return value for environment name, 217
 getexecname — return pathname of executable, 218
 getextmntent — get mnttab file information, 229
 getgrent — group database entry functions, 219
 getgrent_r — group database entry functions, 219
 getgrgid — group database entry functions, 219
 getgrgid_r — group database entry functions, 219
 getgrnam — group database entry functions, 219
 getgrnam_r — group database entry functions, 219
 gethomegroup — obtain information on scheduling decisions, 207
 gethostid — get an identifier for the current host, 223
 gethostname — get name of current host, 224
 gethrtime — get high resolution real time, 225
 gethrvtime — get high resolution virtual time, 225
 getloadavg — get system load averages, 226
 getlogin — get login name, 227
 getlogin_r — get login name, 227
 getmntany — get mnttab file information, 229
 getmntent — get mnttab file information, 229
 getopt — command option parsing, 233
 getopt_clip — parse long command options, 238
 getopt_long — parse long command options, 238
 getopt_long_only — parse long command options, 238
 getpagesize — get system page size, 246

getpass — read a string of characters without echo, 248
 getpassphrase — read a string of characters without echo, 248
 getpeerucred — get connected socket/STREAM peer's credentials, 249
 getpriority — get or set process scheduling priority, 251
 getpw — get passwd entry from UID, 254
 getpwent — get password entry from user database, 255
 getpwent_r — get password entry from user database, 255
 getpwnam — get password entry from user database, 255
 getpwnam_r — get password entry from user database, 255
 getpwuid — get password entry from user database, 255
 getpwuid_r — get password entry from user database, 255
 getrusage — get information about resource utilization, 260
 gets the total number of threads in a process for libc_db — td_ta_get_nthreads, 796
 getspent — get shadow password database entry, 264
 getspent_r — get shadow password database entry (reentrant), 264
 getspnam — get shadow password database entry, 264
 getspnam_r — get shadow password database entry (reentrant), 264
 getsubopt — parse suboptions from a string, 268
 gettext — message handling functions, 272
 gettimeofday — get date and time, 276
 gettimeofday — get system's notion of current Greenwich time, 278
 gettxt — retrieve a text string, 279
 getusershell — get legal user shells, 281
 getutent — user accounting database functions, 282
 getutid — user accounting database functions, 282
 getutline — user accounting database functions, 282
 getutmp — user accounting database functions, 285
 getutmpx — user accounting database functions, 285
 getutxent — user accounting database functions, 285
 getutxid — user accounting database functions, 285
 getutxline — user accounting database functions, 285
 getvfssany — get vfstab file entry, 289
 getvfssent — get vfstab file entry, 289
 getvfssfile — get vfstab file entry, 289
 getvfsspec — get vfstab file entry, 289
 getw — get a byte from a stream, 137
 getwc — get wide character from a stream, 291
 getwchar — get wide character from stdin stream, 292
 getwd — get current working directory pathname, 293
 getwidth — get codeset information, 294
 getws — get a wide-character string from a stream, 295
 getzoneid — map between zone id and name, 297
 getzoneidbyname — map between zone id and name, 297
 getzonenamebyid — map between zone id and name, 297
 glob — generate path names matching a pattern, 298
 globfree — generate path names matching a pattern, 298
 grantpt — grant access to the slave pseudo-terminal device, 302
 group database entry functions — endgrent, 219
 group database entry functions — fgetgrent, 219
 group database entry functions — fgetgrent_r, 219
 group database entry functions — getgrent, 219
 group database entry functions — getgrent_r, 219
 group database entry functions — getgrgid, 219
 group database entry functions — getgrgid_r, 219

- group database entry functions —
getgrnam, 219
- group database entry functions —
getgrnam_r, 219
- group database entry functions — setgrent, 219
- group IDs, supplementary, initialize —
initgroups, 318

H

- halt system processor, — reboot, 628
- hash-table search routine, — hsearch, 303
- hasmntopt — get mnttab file information, 229
- hcreate — create hash table, 303
- hdestroy — destroy hash table, 303
- host name
 - get name of current host —
gethostname, 224
 - set name of current host — sethostname, 224
- hsearch — hash-table search routine, 303

I

- I/O package, standard buffered I/O —
stdio, 711
- iconv — code conversion function, 306
- iconv_close — code conversion deallocation
function, 312
- iconv_open — code conversion allocation
function, 313
- IEEE arithmetic, convert floating-point number
to string — ecvt, 117
- IEEE floating-point environment control
 - fpgetmasks, 161
 - fpgetround, 161
 - fpgetsticky, 161
 - fpsetmask, 161
 - fpsetround, 161
 - fpsetsticky, 161
- imaxabs — return absolute value, 315
- imaxdiv — return quotient and remainder, 316
- index — string operations, 317
- initgroups — initialize the supplementary
group access list, 318
- initialize and destroy mutex attributes object
 - pthread_mutexattr_destroy, 554

- initialize and destroy mutex attributes object
(Continued)
 - pthread_mutexattr_init, 554
- initialize and destroy read-write lock attributes
object
 - pthread_rwlockattr_destroy, 568
 - pthread_rwlockattr_init, 568
- initialize and destroy threads attribute object
 - pthread_attr_destroy, 500
 - pthread_attr_init, 500
- initialize dynamic package —
pthread_once, 566
- initialize or destroy a mutex
 - pthread_mutex_destroy, 559
 - pthread_mutex_init, 559
- initialize or destroy a read-write lock object
 - pthread_rwlock_destroy, 569
 - pthread_rwlock_init, 569
- initialize or destroy condition variable attributes
object
 - pthread_condattr_destroy, 515
 - pthread_condattr_init, 515
- initialize or destroy condition variables
 - pthread_cond_destroy, 517
 - pthread_cond_init, 517
- initstate — pseudorandom number
functions, 610
- input conversion, convert from wchar_t string
— wscanf, 955
- input/output package, standard buffered I/O
— stdio, 711
- insque — insert element to a queue, 319
- interfaces to stdio FILE structure —
_flushlbf, 125
- interfaces to stdio FILE structure —
_fbufsize, 125
- interfaces to stdio FILE structure — __flbf, 125
- interfaces to stdio FILE structure —
__fpending, 125
- interfaces to stdio FILE structure —
__fpurge, 125
- interfaces to stdio FILE structure —
__freadable, 125
- interfaces to stdio FILE structure —
__freading, 125
- interfaces to stdio FILE structure —
__fsetlocking, 125

- interfaces to stdio FILE structure —
 - __fwritable, 125
- interfaces to stdio FILE structure —
 - __fwriting, 125
- invoke isa-specific executable — isaexec, 320
- isaexec — invoke isa-specific executable, 320
- isatty — test for a terminal device, 322
- isnand — determine type of floating-point number, 323
- isnanf — determine type of floating-point number, 323
- iswctype — test character for specified class, 328
- iterate over the set of locks owned by a thread — td_thr_lockowner, 810
- iterator functions on process handles from libc_db — td_ta_sync_iter, 802
- iterator functions on process handles from libc_db — td_ta_thr_iter, 802
- iterator functions on process handles from libc_db — td_ta_tsd_iter, 802

K

- killpg — send signal to a process group, 330

L

- labs — return absolute value of long integer, 26
- language information — nl_langinfo, 413
- ldiv — compute quotient and remainder, 88
- lfind — linear search and update, 346
- lfmt — display error message in standard format and pass to logging and monitoring services, 332
- linear search and update — lfind, 346
- linear search and update — lsearch, 346
- llabs — return absolute value of long long integer, 26
- lldiv — compute quotient and remainder, 88
- lltostr — string conversion routines, 750
- locale, modify and query a program's locale — setlocale, 680
- localeconv — get numeric formatting information, 336

- lock a mutex —
 - pthread_mutex_reltimedlock_np, 564
- lock a mutex — pthread_mutex_timedlock, 564
- lock a read-write lock for reading —
 - pthread_rwlock_reltimedrdlock_np, 573
- lock a read-write lock for reading —
 - pthread_rwlock_timedrdlock, 573
- lock a read-write lock for writing —
 - pthread_rwlock_reltimedwrlock_np, 575
- lock a read-write lock for writing —
 - pthread_rwlock_timedwrlock, 575
- lock a spin lock object —
 - pthread_spin_lock, 593
- lock a spin lock object —
 - pthread_spin_trylock, 593
- lock, apply or remove advisory lock on open file — flock, 145
- lock address space, — mlockall, 386
- lock memory pages, — mlock, 384
- lock or attempt to lock a read-write lock object for reading
 - pthread_rwlock_rdlock, 571
 - pthread_rwlock_tryrdlock, 571
- lock or attempt to lock a read-write lock object for writing
 - pthread_rwlock_trywrlock, 578
 - pthread_rwlock_wrlock, 578
- lock or unlock a mutex
 - pthread_mutex_lock, 561
 - pthread_mutex_trylock, 561
 - pthread_mutex_unlock, 561
- lockf — allows sections of file to be locked, 342
- log message with a stdarg argument list — vsyslog, 898
- login name
 - getlogin, 227
 - getlogin_r, 227
- longjmp — non-local goto, 672, 675
- _longjmp — non-local goto, 345, 675
- lsearch — linear search and update, 346

M

- madvise — provide advice to VM system, 348
- make a FIFO special file — mkfifo, 377
- make a mutex consistent after owner death —
 - pthread_mutex_consistent_np, 555

make a unique file name from a template —
 mktemp, 380
 make a unique file name from a template and
 open the file — mkstemp, 379
 make modified instructions executable —
 sync_instruction_memory, 759
 makecontext — manipulate user contexts, 351
 malloc — memory allocator, 36
 mallocctl — MT hot memory allocator, 392
 manage binary search trees — tdelete, 860
 manage binary search trees — tfind, 860
 manage binary search trees — tsearch, 860
 manage binary search trees — twalk, 860
 manage thread signals for libc_db —
 td_thr_setsigpending, 812
 manage thread signals for libc_db —
 td_thr_sigsetmask, 812
 manage thread-specific data
 — pthread_getspecific, 536
 — pthread_setspecific, 536
 manipulate resource control blocks —
 rctlblk_get_enforced_value, 613
 manipulate resource control blocks —
 rctlblk_get_firing_time, 613
 manipulate resource control blocks —
 rctlblk_get_global_action, 613
 manipulate resource control blocks —
 rctlblk_get_global_flags, 613
 manipulate resource control blocks —
 rctlblk_get_local_action, 613
 manipulate resource control blocks —
 rctlblk_get_local_flags, 613
 manipulate resource control blocks —
 rctlblk_get_privilege, 613
 manipulate resource control blocks —
 rctlblk_get_recipient_pid, 613
 manipulate resource control blocks —
 rctlblk_get_value, 613
 manipulate resource control blocks —
 rctlblk_set_local_action, 613
 manipulate resource control blocks —
 rctlblk_set_local_flags, 613
 manipulate resource control blocks —
 rctlblk_set_privilege, 613
 manipulate resource control blocks —
 rctlblk_set_recipient_pid, 613
 manipulate resource control blocks —
 rctlblk_set_value, 613
 manipulate resource control blocks —
 rctlblk_size, 613
 manipulate sets of signals — sigsetops, 692
 manipulate user contexts — makecontext, 351
 manipulate user contexts — swapcontext, 351
 manipulate sets of signals — sigsetops
 sigaddset, 692
 sigdelset, 692
 sigemptyset, 692
 sigfillset, 692
 sigismember, 692
 map between zone id and name —
 getzoneid, 297
 map between zone id and name —
 getzoneidbyname, 297
 map between zone id and name —
 getzonenamebyid, 297
 match filename or path name — fnmatch, 154
 mblen — get number of bytes in a
 character, 363
 mbrlen — get number of bytes in a character
 (restartable), 364
 mbrtowc — convert a character to a
 wide-character code (restartable), 366
 mbsinit — determine conversion object
 status, 368
 mbsrtowcs — convert a character string to a
 wide-character string (restartable), 369
 mbstowcs — convert a character string to a
 wide-character string, 371
 mbtowc — convert a character to a
 wide-character code, 372
 mctl — memory management control, 373
 memory — memory operations, 375
 memory operations — bcmp, 41
 memory operations — bcopy, 41
 memory operations — bstring, 41
 memory operations — bzero, 41
 memory, optimizing usage of user mapped
 memory — madvise, 348
 memory allocator — bsdmalloc, 36
 memory allocator
 — alloca, 355
 — calloc, 355, 358, 361
 memory allocator — bsdmalloc
 free, 36
 memory allocator
 — free, 355, 358, 361

- memory allocator (Continued)
 - mallinfo, 358
- memory allocator — bsdmalloc
 - malloc, 36
- memory allocator
 - malloc, 355, 358, 361
 - mallopt, 358
 - memalign, 355
- memory allocator — bsdmalloc
 - realloc, 36
- memory allocator
 - realloc, 355, 358, 361
 - valloc, 355
- memory lock or unlock, calling process —
 - plock, 421
- memory management — mctl, 373
- memory management
 - get system page size — getpagesize, 246
 - lock address space — mlockall, 386
 - lock pages in memory — mlock, 384
 - synchronize memory with physical storage
 - msync, 390
 - unlock address space — munlockall, 386
 - unlock pages in memory — munlock, 384
- memory operations
 - memccpy, 375
 - memchr, 375
 - memcmp, 375
 - memcpy, 375
 - memmove, 375
 - memory, 375
 - memset, 375
- message handling functions —
 - bindtextdomain, 272
- message handling functions —
 - bind_textdomain_codeset, 272
- message handling functions — dcgettext, 272
- message handling functions — dcngettext, 272
- message handling functions — dgettext, 272
- message handling functions — dngettext, 272
- message handling functions — gettext, 272
- message handling functions — ngettext, 272
- message handling functions — textdomain, 272
- message catalog
 - open/catalog — catopen, catclose, 44
 - read a program message — catgets, 43

- messages
 - display a message on stderr or system
 - console — fmtmsg, 149
 - print system error messages — perror, 417
 - system signal messages — psignal, 484
- mkfifo — make a FIFO special file, 377
- mkstemp — make a unique file name from a
 - template and open the file, 379
- mktemp — make a unique file name from a
 - template, 380
- mktime — converts a tm structure to a calendar
 - time, 381
- monitor — prepare process execution
 - profile, 388
- msync — synchronize memory with physical
 - storage, 390
- MT hot memory allocator — mallocctl, 392
- MT hot memory allocator — mtmalloc, 392
- mtmalloc — MT hot memory allocator, 392
- mutex_destroy — mutual exclusion locks, 395
- mutex_init — mutual exclusion locks, 395
 - Destroy, 399
 - Dynamically Allocated Mutexes, 405
 - Initialize, 395
 - Interprocess Locking, 402
 - Lock and Unlock, 398
 - Multiple Instruction Single Data, 401
 - Single Gate, 400
 - Solaris Interprocess Robust Locking, 404
- mutex_lock — mutual exclusion locks, 395
- mutex_trylock — mutual exclusion locks, 395
- mutex_unlock — mutual exclusion locks, 395
- mutual exclusion locks
 - mutex_destroy, 395
 - mutex_init, 395
 - mutex_lock, 395
 - mutex_trylock, 395
 - mutex_unlock, 395

N

- ndbm — database functions, 407
- network group entry
 - endnetgrent, 231
 - getnetgrent, 231
 - getnetgrent_r, 231
 - innetgr, 231

- network group entry (Continued)
 - setnetgrent, 231
- nextkey — data base subroutines, 80
- nftw — walk a file tree, 184
- ngettext — message handling functions, 272
- nice — change priority of a process, 411
- nl_langinfo — language information, 413
- nlist — get entries from symbol table, 412
- non-local goto — setjmp, 672, 675
 - _longjmp, 675
 - longjmp, 672, 675
 - _setjmp, 675
 - siglongjmp, 672
 - sigsetjmp, 672
- non-local goto
 - _longjmp, 345
 - _setjmp, 345
- numbers, convert to strings — econvert, 115

O

- obtain information on scheduling decisions —
 - getcpuid, 207
- obtain information on scheduling decisions —
 - gethomedgroup, 207
- offsetof — offset of structure member, 414
- open a shared object — dlopen, 104
- open directory — fdopendir, 415
- open directory — opendir, 415
- open a file — attropen, 34
- open a pseudo terminal device —
 - posix_openpt, 440
- open a stream — fopen, 156
- open a stream — freopen, 171
- opendir — open directory, 415
- openlog — control system log, 769
- operations on a synchronization object in
 - libc_db — td_sync_get_info, 786
- operations on a synchronization object in
 - libc_db — td_sync_get_stats, 786
- operations on a synchronization object in
 - libc_db — td_sync_setstate, 786
- operations on a synchronization object in
 - libc_db — td_sync_waiters, 786
- operations on a synchronization object in
 - libc_db — td_ta_sync_tracking_enable, 786

- output conversion, vsprintf — convert to
 - wchar_t string, 954
- output conversion, formatted
 - fprintf, 471
 - printf, 471
 - sprintf, 471
 - vfprintf, 471
 - vprintf, 471
 - vsprintf, 471

P

- page size, system, get — getpagesize, 246
- parse long command options —
 - getopt_clip, 238
- parse long command options —
 - getopt_long, 238
- parse long command options —
 - getopt_long_only, 238
- password databases
 - lock the lock file — lckpwwdf, 331
 - unlock the lock file — ulckpwwdf, 331
- passwords
 - get password entry from a file —
 - fgetpwent, 255
 - get password entry from a file —
 - fgetpwent_r, 255
 - get password entry in user database —
 - endpwent, 255
 - get password entry in user database —
 - getpwent, 255
 - get password entry in user database —
 - getpwent_r, 255
 - get password entry in user database —
 - getpwnam, 255
 - get password entry in user database —
 - getpwnam_r, 255
 - get password entry in user database —
 - getpwuid, 255
 - get password entry in user database —
 - getpwuid_r, 255
 - get password entry in user database —
 - setpwent, 255
 - get passwd entry from UID — getpw, 254
 - write password file entry — putpwent, 601

passwords, shadow

- get shadow password database entry — endspent, 264
- get shadow password database entry — fgetspent, 264
- get shadow password database entry — getspent, 264
- get shadow password database entry — getspnam, 264
- get shadow password database entry — setspent, 264
- get shadow password database entry (reentrant) — fgetspent_r, 264
- get shadow password database entry (reentrant) — getspent_r, 264
- get shadow password database entry (reentrant) — getspnam_r, 264
- write shadow password file entry — putspent, 603

path name, return last element — path name, 35

pclose — initiate pipe to/from a process, 422

perform word expansions

- wordexp, 950
- wordfree, 950

performs initialization for libc_db library of interfaces — td_init, 784

pperror — print system error messages, 417

pfmt — display error message in standard format, 418

pipes

- initiate to/from a process — pclose, 422
- initiate to/from a process — popen, 422

placeholder for future logging functionality — td_log, 785

plock — lock or unlock into memory process, text, or data, 421

pop a thread cancellation cleanup handler — pthread_cleanup_pop, 510

popen — initiate pipe to/from a process, 422

port_alert — set a port in alert mode, 425

port_associate — associate or dissociate the object with the port, 427

port_create — create a port, 431

port_dissociate — associate or dissociate the object with the port, 427

port_get — retrieve event information from a port, 433

port_getn — retrieve event information from a port, 433

port_send — send a user-defined event to a port or list of ports, 437

port_sendn — send a user-defined event to a port or list of ports, 437

posix_openpt — open a pseudo terminal device, 440

posix_spawn — spawn a process, 442

posix_spawn_file_actions_addclose — add close or open action to spawn file actions object, 456

posix_spawn_file_actions_adddup2 — add dup2 action to spawn file actions object, 458

posix_spawn_file_actions_addopen — add close or open action to spawn file actions object, 456

posix_spawn_file_actions_destroy — destroy and initialize spawn file actions object, 459

posix_spawn_file_actions_init — destroy and initialize spawn file actions object, 459

posix_spawnattr_destroy — destroy and initialize spawn attributes object, 447

posix_spawnattr_getflags — get and set spawn-flags attribute of spawn attributes object, 449

posix_spawnattr_getpgroup — get and set spawn-pgroup attribute of spawn attributes object, 451

posix_spawnattr_getschedparam — get and set spawn-schedparam attribute of spawn attributes object, 452

posix_spawnattr_getschedpolicy — get and set spawn-schedpolicy attribute of spawn attributes object, 453

posix_spawnattr_getsigdefault — get and set spawn-sigdefault attribute of spawn attributes object, 454

posix_spawnattr_getsigmask — get and set spawn-sigmask attribute of spawn attributes object, 455

posix_spawnattr_init — destroy and initialize spawn attributes object, 447

posix_spawnattr_setflags — get and set spawn-flags attribute of spawn attributes object, 449

posix_spawnattr_setpgroup — get and set spawn-pgroup attribute of spawn attributes object, 451
 posix_spawnattr_setschedparam — get and set spawn-schedparam attribute of spawn attributes object, 452
 posix_spawnattr_setschedpolicy — get and set spawn-schedpolicy attribute of spawn attributes object, 453
 posix_spawnattr_setsigdefault — get and set spawn-sigdefault attribute of spawn attributes object, 454
 posix_spawnattr_setsigmask — get and set spawn-sigmask attribute of spawn attributes object, 455
 posix_spawnnp — spawn a process, 442
 preemption control
 — schedctl_exit, 655
 — schedctl_init, 655
 — schedctl_lookup, 655
 — schedctl_start, 655
 — schedctl_stop, 655
 print formatted output
 — fprintf, 460
 — printf, 460
 — snprintf, 460
 — sprintf, 460
 print formatted output of a variable argument list
 — vfprintf, 896
 — vprintf, 896
 — vsnprintf, 896
 — vsprintf, 896
 print formatted wide-character output
 — fwprintf, 190
 — swprintf, 190
 — wprintf, 190
 printf — formatted output conversion, 471
 printf — print formatted output, 460
 printstack — walk stack pointed to by ucontext, 910
 priv_addset — privilege set manipulation functions, 475
 priv_allocset — privilege set manipulation functions, 475
 priv_copyset — privilege set manipulation functions, 475
 priv_delset — privilege set manipulation functions, 475
 priv_emptyset — privilege set manipulation functions, 475
 priv_fillset — privilege set manipulation functions, 475
 priv_freeset — privilege set manipulation functions, 475
 priv_getbyname — privilege name functions, 480
 priv_getbynum — privilege name functions, 480
 priv_getsetbyname — privilege name functions, 480
 priv_getsetbynum — privilege name functions, 480
 priv_gettext — privilege name functions, 480
 priv_ineffect — change privilege sets and check whether privileges are set, 478
 priv_intersect — privilege set manipulation functions, 475
 priv_inverse — privilege set manipulation functions, 475
 priv_isemptyset — privilege set manipulation functions, 475
 priv_isequalset — privilege set manipulation functions, 475
 priv_isfullset — privilege set manipulation functions, 475
 priv_ismember — privilege set manipulation functions, 475
 priv_issubset — privilege set manipulation functions, 475
 priv_set — change privilege sets and check whether privileges are set, 478
 priv_set_to_str — privilege name functions, 480
 priv_str_to_set — privilege name functions, 480
 priv_union — privilege set manipulation functions, 475
 privilege name functions —
 priv_getbyname, 480
 priv_getbynum, 480
 priv_getsetbyname, 480

- privilege name functions —
 - priv_getsetbynum, 480
- privilege name functions — priv_gettext, 480
- privilege name functions —
 - priv_set_to_str, 480
- privilege name functions —
 - priv_str_to_set, 480
- privilege set manipulation functions —
 - priv_addset, 475
- privilege set manipulation functions —
 - priv_allocset, 475
- privilege set manipulation functions —
 - priv_copysset, 475
- privilege set manipulation functions —
 - priv_delset, 475
- privilege set manipulation functions —
 - priv_emptyset, 475
- privilege set manipulation functions —
 - priv_fillset, 475
- privilege set manipulation functions —
 - priv_freerset, 475
- privilege set manipulation functions —
 - priv_intersect, 475
- privilege set manipulation functions —
 - priv_inverse, 475
- privilege set manipulation functions —
 - priv_isemptyset, 475
- privilege set manipulation functions —
 - priv_isequalset, 475
- privilege set manipulation functions —
 - priv_isfullset, 475
- privilege set manipulation functions —
 - priv_ismember, 475
- privilege set manipulation functions —
 - priv_issubset, 475
- privilege set manipulation functions —
 - priv_union, 475
- Process Code string operations — wstring, 956
- Process Code string operations — wscasecmp, wscasecmp, 956
- Process Code string operations — wscol, wscol, 956
- Process Code string operations — wsdup, wsdup, 956
- Process Code string operations — wscasecmp, wscasecmp, 956
- process statistics, prepare execution profile —
 - monitor, 388
- processes
 - allows a parent process to control the
 - execution of a child process — ptrace, 596
 - change priority — nice, 411
 - duplicate an open file descriptor —
 - dup2, 114
 - generate path name for controlling terminal
 - ctermid, ctermid_r, 70
 - get character-string representation —
 - cuserid, 79
 - initiate pipe to/from a process — popen, pclose, 422
 - memory lock or unlock — plock, 421
 - prepare execution profile — monitor, 388
 - report CPU time used — clock, 49
 - send signal to a process group — killpg, 330
 - send signal to program — raise, 607
 - suspend execution for interval — sleep, 703
 - terminate process — exit, 122
 - terminate the process abnormally —
 - abort, 25
 - wait for process to terminate or stop —
 - WIFEXITED, 904
 - wait for process to terminate or stop —
 - WIFSIGNALED, 904
 - wait for process to terminate or stop —
 - WIFSTOPPED, 904
 - wait for process to terminate or stop —
 - wait, 904
 - wait for child process to change state —
 - waitpid, 908
 - wait for child process to stop or terminate —
 - wait, 902
- profiling utilities, prepare process execution
 - profile — monitor, 388
- program assertion, verify — assert, 30
- program messages
 - open/close a message catalog — catopen, catclose, 44
 - read — catgets, 43
- programs, last locations — end, etext, edata, 120
- pselect — synchronous I/O multiplexing, 658
- pset_getloadavg — get system load averages for a processor set, 483
- pseudo-terminal device
 - get name of the slave pseudo-terminal device
 - ptsname, 599

pseudo-terminal device (Continued)
 grant access to the slave pseudo-terminal device — `grantpt`, 302
 pseudorandom number functions
 — `initstate`, 610
 — `random`, 610
 — `setstate`, 610
 — `srandom`, 610
`psiginfo` — system signal messages, 484
`psignal` — system signal messages, 484, 485
`pthread_atfork` — register fork handlers, 486
`pthread_attr_destroy` — initialize and destroy threads attribute object, 500
`pthread_attr_getdetachstate` — get or set detachstate attribute, 488
`pthread_attr_getguardsize` — get or set the thread guardsize attribute, 489
`pthread_attr_getinheritsched` — get or set inheritsched attribute, 491
`pthread_attr_getschedparam` — get or set schedparam attribute, 493
`pthread_attr_getschedpolicy` — get or set schedpolicy attribute, 494
`pthread_attr_getscope` — get or set contentionscope attribute, 495
`pthread_attr_getstack` — get and set stack attributes, 496
`pthread_attr_getstackaddr` — get or set stackaddr attribute, 498
`pthread_attr_getstacksize` — get or set stacksize attribute, 499
`pthread_attr_init` — initialize and destroy threads attribute object, 500
`pthread_attr_setdetachstate` — get or set detachstate attribute, 488
`pthread_attr_setguardsize` — get or set the thread guardsize attribute, 489
`pthread_attr_setinheritsched` — get or set inheritsched attribute, 491
`pthread_attr_setschedparam` — get or set schedparam attribute, 493
`pthread_attr_setschedpolicy` — get or set schedpolicy attribute, 494
`pthread_attr_setscope` — get or set contentionscope attribute, 495
`pthread_attr_setstack` — get and set stack attributes, 496
`pthread_attr_setstackaddr` — get or set stackaddr attribute, 498
`pthread_attr_setstacksize` — get or set stacksize attribute, 499
`pthread_barrier_destroy` — destroy and initialize a barrier object, 505
`pthread_barrier_init` — destroy and initialize a barrier object, 505
`pthread_barrier_wait` — synchronize at a barrier, 507
`pthread_barrierattr_destroy` — destroy and initialize barrier attributes object, 502
`pthread_barrierattr_getpshared` — get and set process-shared attribute of barrier attributes object, 503
`pthread_barrierattr_init` — destroy and initialize barrier attributes object, 502
`pthread_barrierattr_setpshared` — get and set process-shared attribute of barrier attributes object, 503
`pthread_cleanup_pop` — pop a thread cancellation cleanup handler, 510
`pthread_cleanup_push` — push a thread cancellation cleanup handler, 511
`pthread_cond_broadcast` — signal or broadcast a condition, 519
`pthread_cond_destroy` — initialize or destroy condition variables, 517
`pthread_cond_init` — initialize or destroy condition variables, 517
`pthread_cond_reltimedwait_np` — wait on a condition, 521
`pthread_cond_signal` — signal or broadcast a condition, 519
`pthread_cond_timedwait` — wait on a condition, 521
`pthread_cond_wait` — wait on a condition, 521
`pthread_condattr_destroy` — initialize or destroy condition variable attributes object, 515
`pthread_condattr_getclock` — get and set the clock selection condition variable attribute, 512
`pthread_condattr_getpshared` — get or set the process-shared condition variable attributes, 513
`pthread_condattr_init` — initialize or destroy condition variable attributes object, 515

`pthread_condattr_setclock` — get and set the clock selection condition variable attribute, 512
`pthread_condattr_setpshared` — get or set the process-shared condition variable attributes, 513
`pthread_create` — create a thread, 525
`pthread_detach` — detach a thread, 529
`pthread_equal` — compare thread IDs, 530
`pthread_exit` — terminate calling thread, 531
`pthread_getconcurrency` — get or set level of concurrency, 532
`pthread_getschedparam` — access dynamic thread scheduling parameters, 534
`pthread_getspecific` — manage thread-specific data, 536
`pthread_join` — wait for thread termination, 537
`pthread_key_create` — create thread-specific data key, 539
`pthread_key_delete` — delete thread-specific data key, 541
`pthread_mutex_consistent_np` — make a mutex consistent after owner death, 555
`pthread_mutex_destroy` — initialize or destroy a mutex, 559
`pthread_mutex_getprioceiling` — change the priority ceiling of a mutex, 557
`pthread_mutex_init` — initialize or destroy a mutex, 559
`pthread_mutex_lock` — lock or unlock a mutex, 561
`pthread_mutex_reltimedlock_np` — lock a mutex, 564
`pthread_mutex_setprioceiling` — change the priority ceiling of a mutex, 557
`pthread_mutex_timedlock` — lock a mutex, 564
`pthread_mutex_trylock` — lock or unlock a mutex, 561
`pthread_mutex_unlock` — lock or unlock a mutex, 561
`pthread_mutexattr_destroy` — initialize and destroy mutex attributes object, 554
`pthread_mutexattr_getprioceiling` — get and set prioceiling attribute of mutex attribute object, 543
`pthread_mutexattr_getprotocol` — get and set protocol attribute of mutex attribute object, 545
`pthread_mutexattr_getpshared` — get and set process-shared attribute, 548
`pthread_mutexattr_gettype` — get or set a mutex type, 552
`pthread_mutexattr_init` — initialize and destroy mutex attributes object, 554
`pthread_mutexattr_setprioceiling` — get and set prioceiling attribute of mutex attribute object, 543
`pthread_mutexattr_setprotocol` — get and set protocol attribute of mutex attribute object, 545
`pthread_mutexattr_setpshared` — get and set process-shared attribute, 548
`pthread_mutexattr_settype` — get or set a mutex type, 552
`pthread_once` — initialize dynamic package, 566
`pthread_rwlock_destroy` — initialize or destroy a read-write lock object, 569
`pthread_rwlock_init` — initialize or destroy a read-write lock object, 569
`pthread_rwlock_rdlock` — lock or attempt to lock a read-write lock object for reading, 571
`pthread_rwlock_reltimedrdlock_np` — lock a read-write lock for reading, 573
`pthread_rwlock_reltimedwrlock_np` — lock a read-write lock for writing, 575
`pthread_rwlock_timedrdlock` — lock a read-write lock for reading, 573
`pthread_rwlock_timedwrlock` — lock a read-write lock for writing, 575
`pthread_rwlock_tryrdlock` — lock or attempt to lock a read-write lock object for reading, 571
`pthread_rwlock_trywrlock` — lock or attempt to lock a read-write lock object for writing, 578
`pthread_rwlock_unlock` — unlock a read-write lock object, 577
`pthread_rwlock_wrlock` — lock or attempt to lock a read-write lock object for writing, 578
`pthread_rwlockattr_destroy` — initialize and destroy read-write lock attributes object, 568
`pthread_rwlockattr_getpshared` — get and set process-shared attribute of read-write lock attributes object, 567

pthread_rwlockattr_init — initialize and destroy read-write lock attributes object, 568
 pthread_rwlockattr_setpshared — get and set process-shared attribute of read-write lock attributes object, 567
 pthread_self — get calling thread's ID, 580
 pthread_setcancelstate — enable or disable cancellation, 581
 pthread_setcancelstate — enabling or disabling cancellation, 581
 pthread_setcanceltype — set the cancellation type of a thread, 583
 pthread_setconcurrency — get or set level of concurrency, 532
 pthread_setschedparam — access dynamic thread scheduling parameters, 534
 pthread_setschedprio — dynamic thread scheduling parameters access, 585
 pthread_setspecific — manage thread-specific data, 536
 pthread_sigmask — change or examine calling thread's signal mask, 586
 pthread_spin_destroy — destroy or initialize a spin lock object, 591
 pthread_spin_init — destroy or initialize a spin lock object, 591
 pthread_spin_lock — lock a spin lock object, 593
 pthread_spin_trylock — lock a spin lock object, 593
 pthread_spin_unlock — unlock a spin lock object, 594
 pthread_testcancel — create cancellation point in the calling thread., 595
 ptrace — allows a parent process to control the execution of a child process, 596
 ptsname — get name of the slave pseudo-terminal device, 599
 push a thread cancellation cleanup handler — pthread_cleanup_push, 511
 push byte back into input stream — ungetc, 886
 push wide-character code back into input stream — ungetwc, 887
 put wide-character code on a stream — fputwc, 166
 put wide-character code on a stream — putwc, 166
 put wide-character code on a stream — putwchar, 166
 put a byte on a stream — fputc, 163
 — putc, 163
 — putc_unlocked, 163
 — putchar, 163
 — putchar_unlocked, 163
 — putw, 163
 put wide character string on a stream — fputws, 168
 putc — put a byte on a stream, 163
 putc_unlocked — put a byte on a stream, 163
 putchar — put a byte on a stream, 163
 putchar_unlocked — put a byte on a stream, 163
 putenv — change or add value to environment, 600
 putmntent — get mnttab file information, 229
 putpwent — write password file entry, 601
 putsptent — write shadow password file entry, 603
 pututline — user accounting database functions, 282
 pututxline — user accounting database functions, 285
 putw — put a byte on a stream, 163
 putwc — put wide-character code on a stream, 166
 putwchar — put wide-character code on a stream, 166
 putws — convert a string of Process Code characters to EUC characters and put it on a stream, 604

Q

qeconvert — convert number to ASCII, 115
 qfconvert — convert number to ASCII, 115
 qgconvert — convert number to ASCII, 115
 qsort — quick sort, 605
 quadruple_to_decimal — decimal record from quadruple-precision floating, 143
 queues, insert/remove element from a queue — insque, remque, 319
 quick sort — qsort, 605

R

- raise — send signal to program, 607
- rand — simple random number generator, 609
- rand — simple random-number generator, 608
- random — pseudorandom number functions, 610
- random number generator
 - drand48, 112
 - erand48, 112
 - jrand48, 112
 - lcong48, 112
 - lrand48, 112
 - mrand48, 112
 - nrand48, 112
 - rand, 609
 - seed48, 112
 - srand48, 112
- random number generator, simple
 - rand, 608
 - srand, 608
- rctl_walk — visit registered rctls on current system, 618
- rctlblk_get_enforced_value — manipulate resource control blocks, 613
- rctlblk_get_firing_time — manipulate resource control blocks, 613
- rctlblk_get_global_action — manipulate resource control blocks, 613
- rctlblk_get_global_flags — manipulate resource control blocks, 613
- rctlblk_get_local_action — manipulate resource control blocks, 613
- rctlblk_get_local_flags — manipulate resource control blocks, 613
- rctlblk_get_privilege — manipulate resource control blocks, 613
- rctlblk_get_recipient_pid — manipulate resource control blocks, 613
- rctlblk_get_value — manipulate resource control blocks, 613
- rctlblk_set_local_action — manipulate resource control blocks, 613
- rctlblk_set_local_flags — manipulate resource control blocks, 613
- rctlblk_set_privilege — manipulate resource control blocks, 613
- rctlblk_set_recipient_pid — manipulate resource control blocks, 613
- rctlblk_set_value — manipulate resource control blocks, 613
- rctlblk_size — manipulate resource control blocks, 613
- re_comp — compile and execute regular expressions, 629
- re_exec — compile and execute regular expressions, 629
- read a string of characters without echo — getpass, 248
- read a string of characters without echo — getpassphrase, 248
- read directory — readdir, 620
- read directory — readdir_r, 620
- read a directory entry — readdir, 624
- read directory
 - readdir, 620
 - readdir_r, 620
- readdir — read a directory entry, 624
- readdir — read directory, 620
- readdir_r — read directory, 620
- reading and writing thread registers in libc_db — td_thr_getxregs, 805
- reading and writing thread registers in libc_db — td_thr_getfpregs, 805
- reading and writing thread registers in libc_db — td_thr_getgregs, 805
- reading and writing thread registers in libc_db — td_thr_getxregsize, 805
- reading and writing thread registers in libc_db — td_thr_setfpregs, 805
- reading and writing thread registers in libc_db — td_thr_setgregs, 805
- reading and writing thread registers in libc_db — td_thr_setxregs, 805
- realloc — memory allocator, 36
- realpath — resolve pathname, 626
- reboot — reboot system or halt processor, 628
- regcmp — compile regular expression, 630
- regcomp — regular expression matching, 632
- regerror — regular expression matching, 632
- regex — execute regular expression, 630
- regexec — regular expression matching, 632
- regfree — regular expression matching, 632
- register a function to run at process termination or object unloading— atexit, 31
- register fork handlers — pthread_atfork, 486

regular expression matching
 — regcomp, 632
 — regerror, 632
 — regexec, 632
 — regfree, 632
 regular expressions, compile and execute —
 regcmp, regex, 630
 remove — remove file, 638
 remove an environment variable —
 unsetenv, 889
 remove — remove element from a queue, 319
 reposition a file-position indicator in a stream
 — fseek, 174
 — fseeko, 174
 reset file position indicator in a stream —
 rewind, 639
 reset position of directory stream to the
 beginning of a directory — rewinddir, 640
 resetmnttab — get mnttab file information, 229
 resolve pathname — realpath, 626
 resource utilization, get information —
 getrusage, 260
 retrieve event information from a port —
 port_get, 433
 retrieve event information from a port —
 port_getn, 433
 retrieve stack boundaries —
 stack_getbounds, 705
 return a file offset for a file descriptor —
 tell, 816
 return a file offset in a stream
 — ftell, 180
 — ftello, 180
 return absolute value — imaxabs, 315
 return pathname of executable —
 getexecname, 218
 return quotient and remainder — imaxdiv, 316
 return the synchronization handle for the object
 on which a thread is blocked —
 td_thr_sleepinfo, 813
 rewind — reset file position indicator in a
 stream, 639
 rewinddir — reset position of directory stream
 to the beginning of a directory, 640
 rindex — string operations, 317
 rw_rdlock() — acquire a read lock, 641
 rw_tryrdlock() — acquire a read lock, 641
 rw_trywrlock() — acquire a write lock, 641

rw_unlock() — unlock a readers/writer
 lock, 641
 rw_wrlock() — acquire a write lock, 641
 rwlock_destroy() — destroy a readers/writer
 lock, 641
 rwlock_init() — initialize a readers/writer
 lock, 641

S

scan a directory — alphasort, 644
 scan a directory — scandir, 644
 scan a directory
 — alphasort, 646
 — scandir, 646
 scandir — scan a directory, 644, 646
 scanf — convert formatted input, 647
 Conversion Characters, 650
 Conversion Specifications, 648
 schedctl_exit — preemption control, 655
 schedctl_init — preemption control, 655
 schedctl_lookup — preemption control, 655
 schedctl_start — preemption control, 655
 schedctl_stop — preemption control, 655
 scheduling priority, change priority of a process
 — nice, 411
 search functions, manage hash search tables —
 hsearch, 303
 seconvert — convert number to ASCII, 115
 seekdir — set position of directory stream, 657
 select — synchronous I/O multiplexing, 658
 sema_destroy() — destroy a semaphore, 663
 sema_init() — initialize a semaphore, 663
 sema_post() — increment a semaphore, 663
 sema_trywait() — decrement a semaphore, 663
 sema_wait() — decrement a semaphore, 663
 send a user-defined event to a port or list of
 ports — port_send, 437
 send a user-defined event to a port or list of
 ports — port_sendn, 437
 send a “break” for a specific duration —
 tcseendbreak, 780
 set a port in alert mode — port_alert, 425
 set and/or get alternate signal stack context —
 sigstack, 694
 set and/or get signal stack context —
 sigstack, 696

set concurrency level for target process —

 set encoding key — setkey, 678
 set foreground process group ID —
 tcsetpgrp, 783
 set input baud rate, — cfsetispeed, 48
 set output baud rate, — cfsetospeed, 48
 set position of directory stream — seekdir, 657
 set stream orientation — fwide, 189
 set the cancellation type of a thread —
 pthread_setcanceltype, 583
 set the parameters associated with the terminal
 — tcsetattr, 781
 set the priority of a thread —
 td_thr_setprio, 811
 set wide-characters in memory —
 wmemset, 949
 setcat — define default catalog, 670
 setenv — add or change environment
 variable, 671
 setgrent — group database entry functions, 219
 sethostname — set name of current host, 224
 setjmp — non-local goto, 672, 675
 _setjmp — non-local goto, 345, 675
 setkey — set encoding key, 678
 setlabel — define the label for pfmt() and
 lfmt()., 679
 setlocale — modify and query a program's
 locale, 680
 setlogmask — control system log, 769
 setpriority — get or set process scheduling
 priority, 251
 setpwnam — get password entry from user
 database, 255
 setspent — get shadow password database
 entry, 264
 setstate — pseudorandom number
 functions, 610
 settimeofday — set date and time, 276
 settimeofday — set system's notion of current
 Greenwich time, 278
 setusershell — get legal user shells, 281
 setutent — user accounting database
 functions, 282
 setutxent — user accounting database
 functions, 285
 severity levels, applications, build a list for use
 with fmtmsg — addseverity, 28

 sfcvrt — convert number to ASCII, 115
 sgconvrt — convert number to ASCII, 115
 shared object
 close — dlclose, 91
 get address of symbol — dlsym, 109
 get diagnostic information — dlerror, 99
 open — dlopen, 104
 shell command, issue one — system, 773
 sig2str — translation between signal name and
 signal number, 715
 sigaddset — manipulate sets of signals, 692
 sigdelset — manipulate sets of signals, 692
 sigemptyset — manipulate sets of signals, 692
 sigfillset — manipulate sets of signals, 692
 sigfpe() function, 684
 sighold — adds sig to the calling process's
 signal mask, 688
 sigignore — sets the disposition of sig to
 SIG_IGN, 688
 siginterrupt — allow signals to interrupt
 functions, 686, 687
 sigismember — manipulate sets of signals, 692
 siglongjmp — non-local goto, 672
 signal — modify signal disposition, 688
 signal — simplified software signal
 facilities, 690
 signal
 schedule after interval in microseconds —
 ualarm, 866
 suspend execution for interval in
 microseconds — usleep, 890
 simplified signal facilities — bsd_signal, 38
 signal management, simplified, for application
 processes — signal, 688
 signal messages, system
 — psignal, 484, 485
 signal or broadcast a condition
 — pthread_cond_broadcast, 519
 — pthread_cond_signal, 519
 signals, block
 — sigblock, 683
 — sigmask, 683
 — sigpause, 683
 — sigsetmask, 683
 signals, software
 — gsignal, 704
 — ssignal, 704

sigpause — removes sig from the calling process's signal mask and suspends the calling process until a signal is received, 688
 sigrelse — removes sig from the calling process's signal mask, 688
 sigset — modify signal disposition, 688
 sigsetjmp — non-local goto, 672
 sigsetops — manipulate sets of signals, 692
 sigstack — set and/or get alternate signal stack context, 694
 sigstack — set and/or get signal stack context, 696
 sigvec — software signal facilities, 697
 single-byte to wide-character conversion — btowc, 42
 single_to_decimal — decimal record from single-precision floating, 143
 sleep — suspend execution for interval, 703
 sleep, suspend execution for interval in microseconds — usleep, 890
 snprintf — print formatted output, 460
 software signals
 — gsignal, 704
 — ssignal, 704
 spawn a process — posix_spawn, 442
 spawn a process — posix_spawn, 442
 sprintf — formatted output conversion, 471
 printf — print formatted output, 460
 srand — reset simple random number generator, 609
 random — pseudorandom number functions, 610
 sscanf — convert formatted input, 647
 stack_getbounds — retrieve stack boundaries, 705
 _stack_grow — express an intention to extend the stack, 706
 stack_inbounds — determine if address is within stack boundaries, 707
 stack_setbounds — update stack boundaries, 708
 stack_violation — determine stack boundary violation event, 709
 stdio — standard buffered input/output package, 711
 sting collation, — strcoll, 716
 store — data base subroutines, 80
 str2sig — translation between signal name and signal number, 715
 strcasecmp — string operations, 728
 strcat — string operations, 728
 strchr — string operations, 728
 strcmp — string operations, 728
 strcpy — string operations, 728
 strcspn — string operations, 728
 strdup — string operations, 728
 stream
 convert a string of Process Code characters to EUC characters and put it on a stream — putws, 604
 open — fopen, 159
 stream, assign buffering
 — setbuf, 667
 — setvbuf, 667
 stream, get string
 — fgets, 263
 — gets, 263
 stream, put a string
 — fputs, 602
 — puts, 602
 stream status inquiries
 — clearerr, 133
 — feof, 133
 — ferror, 133
 — fileno, 133
 STREAMS, test file descriptor for a STREAMS file — isastream, 321
 strerror — get error message string, 717
 strerror_r — get error message string, 717
 strfmon — convert monetary value to string, 718
 strftime — convert date and time to string, 722
 string — string operations, 728
 string operations — strcasecmp, 728
 string operations — strcat, 728
 string operations — strchr, 728
 string operations — strcmp, 728
 string operations — strcpy, 728
 string operations — strcspn, 728
 string operations — strdup, 728
 string operations — string, 728
 string operations — strlcat, 728
 string operations — strlcpy, 728
 string operations — strlen, 728
 string operations — strncasecmp, 728

string operations — `strncat`, 728
 string operations — `strncmp`, 728
 string operations — `strncpy`, 728
 string operations — `strpbrk`, 728
 string operations — `strchr`, 728
 string operations — `strspn`, 728
 string operations — `strstr`, 728
 string operations — `strtok`, 728
 string operations — `strtok_r`, 728
 string conversion routines
 — `atoi`, 750
 — `atol`, 750
 — `atoll`, 750
 — `lltostr`, 750
 — `strtol`, 750
 — `strtoll`, 750
 — `ulltostr`, 750
 string encoding function — `crypt`, 64
 string operations
 — `index`, 317
 — `rindex`, 317
 string_to_decimal — decimal record from character string, 733
 string transformation — `strxfrm`, 756
 strings, convert from numbers — `econvert`, 115
`strlcat` — string operations, 728
`strlcpy` — string operations, 728
`strlen` — string operations, 728
`strncasecmp` — string operations, 728
`strncat` — string operations, 728
`strncmp` — string operations, 728
`strncpy` — string operations, 728
`strpbrk` — string operations, 728
`strptime` — date and time conversion, 737
`strchr` — string operations, 728
`strsignal` — get name of signal, 743
`strspn` — string operations, 728
`strstr` — string operations, 728
`strtod` — convert string to floating-point number, 744
`strtof` — convert string to floating-point number, 744
`strtok` — string operations, 728
`strtok_r` — string operations, 728
`strtol` — string conversion routines, 750
`strtold` — convert string to floating-point number, 744
`strtoll` — string conversion routines, 750
`strtoul` — convert string to unsigned long, 753
`strtows` — code conversion for Process Code and File Code, 755
`strxfrm` — string transformation, 756
 suspend and resume threads in `libc_db` — `td_thr_dbresume`, 804
 suspend and resume threads in `libc_db` — `td_thr_dbsuspend`, 804
 suspend or restart the transmission or reception of data — `tcflow`, 775
`swab` — swap bytes, 758
 swap bytes — `swab`, 758
`swapcontext` — manipulate user contexts, 351
`swprintf` — print formatted wide-character output, 190
`swscanf` — convert formatted wide-character input, 200
 symbol address, get address in shared object or executable — `dlsym`, 109
 symbol table, get entries — `nlist`, 412
`sync_instruction_memory` — make modified instructions executable, 759
 synchronize at a barrier — `pthread_barrier_wait`, 507
 synchronize changes to a file — `fsync`, 178
 synchronous I/O multiplexing — `FD_CLR`, 658
 synchronous I/O multiplexing — `FD_ISSET`, 658
 synchronous I/O multiplexing — `FD_SET`, 658
 synchronous I/O multiplexing — `FD_ZERO`, 658
 synchronous I/O multiplexing — `pselect`, 658
 synchronous I/O multiplexing — `select`, 658
`sys_siglist` — system signal messages list, 485
`syscall` — indirect system call, 760
`sysconf` — get configurable system variables, 761
`syslog` — control system log, 769
 system — issue shell command, 773
 system error messages, print — `perror`, 417
 system signal messages, — `psignal`, 484
 system variables, get configurable ones — `sysconf`, 761

T
`tcdrain` — wait for transmission of output, 774

`tcflow` — suspend or restart the transmission or reception of data, 775
`tcflush` — flush non-transmitted output data, non-read input data or both, 776
`tcgetattr` — get the parameters associated with the terminal, 777
`tcgetpgrp` — get foreground process group ID, 778
`tcgetsid` — get process group ID for session leader for controlling terminal, 779
`tcsendbreak` — send a “break” for a specific duration, 780
`tcsetattr` — set the parameters associated with the terminal, 781
`tcsetpgrp` — set foreground process group ID, 783
`td_event_addset` — thread events in `libc_db`, 792
`td_event_delset` — thread events in `libc_db`, 792
`td_event_emptyset` — thread events in `libc_db`, 792
`td_event_fillset` — thread events in `libc_db`, 792
`td_eventisempty` — thread events in `libc_db`, 792
`td_eventismember` — thread events in `libc_db`, 792
`td_init` — performs initialization for `libc_db` library of interfaces, 784
`td_log` — placeholder for future logging functionality, 785
`td_sync_get_info` — operations on a synchronization object in `libc_db`, 786
`td_sync_get_stats` — operations on a synchronization object in `libc_db`, 786
`td_sync_setstate` — operations on a synchronization object in `libc_db`, 786
`td_sync_waiters` — operations on a synchronization object in `libc_db`, 786
`td_ta_clear_event` — thread events in `libc_db`, 792
`td_ta_delete` — allocate and deallocate process handles for `libc_db`, 799
`td_ta_enable_stats` — collect target process statistics for `libc_db`, 790
`td_ta_event_addr` — thread events in `libc_db`, 792
`td_ta_event_getmsg` — thread events in `libc_db`, 792
`td_ta_get_nthreads` — gets the total number of threads in a process for `libc_db`, 796
`td_ta_get_ph` — allocate and deallocate process handles for `libc_db`, 799
`td_ta_get_stats` — collect target process statistics for `libc_db`, 790
`td_ta_map_addr2sync` — get a synchronization object handle from a synchronization object’s address, 797
`td_ta_map_id2thr` — convert a thread ID or LWP ID to a thread handle, 798
`td_ta_map_lwp2thr` — convert a thread ID or LWP ID to a thread handle, 798
`td_ta_new` — allocate and deallocate process handles for `libc_db`, 799
`td_ta_reset_stats` — collect target process statistics for `libc_db`, 790
`td_ta_set_event` — thread events in `libc_db`, 792
`td_ta_setconcurrency` — set concurrency level for target process, 801
`td_ta_sync_iter` — iterator functions on process handles from `libc_db`, 802
`td_ta_sync_tracking_enable` — operations on a synchronization object in `libc_db`, 786
`td_ta_thr_iter` — iterator functions on process handles from `libc_db`, 802
`td_ta_tsd_iter` — iterator functions on process handles from `libc_db`, 802
`td_thr_clear_event` — thread events in `libc_db`, 792
`td_thr_dbresume` — suspend and resume threads in `libc_db`, 804
`td_thr_dbsuspend` — suspend and resume threads in `libc_db`, 804
`td_thr_event_enable` — thread events in `libc_db`, 792
`td_thr_event_getmsg` — thread events in `libc_db`, 792
`td_thr_get_info` — get thread information in `libc_db` library of interfaces, 807
`td_thr_getfpregs` — reading and writing thread registers in `libc_db`, 805
`td_thr_getgregs` — reading and writing thread registers in `libc_db`, 805

td_thr_getxregs — reading and writing thread registers in libc_db, 805
 td_thr_getxregsize — reading and writing thread registers in libc_db, 805
 td_thr_lockowner — iterate over the set of locks owned by a thread, 810
 td_thr_set_event — thread events in libc_db, 792
 td_thr_setfpregs — reading and writing thread registers in libc_db, 805
 td_thr_setgregs — reading and writing thread registers in libc_db, 805
 td_thr_setprio — set the priority of a thread, 811
 td_thr_setsigpending — manage thread signals for libc_db, 812
 td_thr_setxregs — reading and writing thread registers in libc_db, 805
 td_thr_sigsetmask — manage thread signals for libc_db, 812
 td_thr_sleepinfo — return the synchronization handle for the object on which a thread is blocked, 813
 td_thr_tsd — get a thread's thread-specific data for libc_db library of interfaces, 814
 td_thr_validate — test a thread handle for validity, 815
 tdelete — manage binary search trees, 860
 tell — return a file offset for a file descriptor, 816
 telldir — current location of a named directory stream, 817
 tempnam — create a name for a temporary file, 847
 terminal, find the slot of the current user in the user accounting database — ttyslot, 865
 terminal device, slave pseudo
 get name — ptsname, 599
 grant access — grantpt, 302
 terminal ID, generate path name for controlling terminal — ctermid, ctermid_r, 70
 terminate calling thread — pthread_exit, 531
 terminate the calling thread — thr_exit, 825
 termios — general terminal interface, 818
 test a thread handle for validity — td_thr_validate, 815
 test character for specified class — iswctype, 328
 test for a terminal device — isatty, 322
 text processing utilities, compile and execute regular expressions — regcmp, regex, 630
 text string, — gettxt, 279
 textdomain — message handling functions, 272
 tfind — manage binary search trees, 860
 thr_continue — continue thread execution, 843
 thr_create — create a thread, 819
 thr_exit — terminate the calling thread, 825
 thr_getconcurrency — get thread concurrency level, 827
 thr_getprio — access dynamic thread scheduling, 828
 thr_getspecific — thread-specific-data functions, 831
 thr_join — wait for thread termination, 829
 thr_keycreate — thread-specific-data functions, 831
 Create Key, 831
 Get Value, 831
 Set Value, 831
 thr_main — identifies the calling thread as the main thread or not the main thread, 835
 thr_self — get calling thread's ID, 837
 thr_setconcurrency — set thread concurrency level, 827
 thr_setprio — access dynamic thread scheduling, 828
 thr_setspecific — thread-specific-data functions, 831
 thr_sigsetmask — change or examine calling thread's signal mask, 838
 thr_stksegment — get thread stack address and size, 842
 thr_suspend — suspend thread execution, 843
 thr_yield — thread yield to another thread, 844
 thread events in libc_db —
 td_eventisempty, 792
 thread events in libc_db —
 td_eventismember, 792
 thread events in libc_db —
 td_event_addset, 792
 thread events in libc_db —
 td_event_delset, 792
 thread events in libc_db —
 td_event_emptyset, 792
 thread events in libc_db —
 td_event_fillset, 792

thread events in libc_db —
 td_ta_clear_event, 792
 thread events in libc_db —
 td_ta_event_addr, 792
 thread events in libc_db —
 td_ta_event_getmsg, 792
 thread events in libc_db —
 td_ta_set_event, 792
 thread events in libc_db —
 td_thr_clear_event, 792
 thread events in libc_db —
 td_thr_event_enable, 792
 thread events in libc_db —
 td_thr_event_getmsg, 792
 thread events in libc_db —
 td_thr_set_event, 792
 thread-specific-data functions —
 thr_getspecific, 831
 thread-specific-data functions —
 thr_keycreate, 831
 thread-specific-data functions —
 thr_setspecific, 831
 thread yield to another thread — thr_yield, 844
 time, computes the difference between two
 calendar times — difftime, 83
 time, calendar, convert from a tm structure —
 mktime, 381
 time accounting, for current process —
 times, 845
 time and date
 convert to string — asctime, 71
 convert to string — ctime, 71
 convert to string — gmtime, 71
 convert to string — localtime, 71
 convert to string — tzset, 71
 convert user format date and time —
 getdate, 210
 get — ftime, 181
 — settimeofday, 276
 time of day, get and set — gettimeofday,
 settimeofday, 278
 times — get process times, 845
 tmpfile — create a temporary file, 846
 tmpnam — create a name for a temporary
 file, 847
 toascii — translate integer to a 7-bit ASCII
 character, 849
 tolower — transliterate upper-case characters to
 lower-case, 851
 _tolower — transliterate upper-case characters
 to lower-case, 850
 toupper — transliterate lower-case characters to
 upper-case, 853
 _toupper — transliterate lower-case characters
 to upper-case, 852
 towctrans — wide-character mapping, 854
 tolower — transliterate upper-case
 wide-character code to lower-case, 855
 toupper — transliterate lower-case
 wide-character code to upper-case, 856
 translate address to symbolic information —
 dladdr1, 89
 translate address to symbolic information —
 dladdr, 89
 translate integer to a 7-bit ASCII character —
 toascii, 849
 translation between signal name and signal
 number — str2sig, 715
 sig2str, 715
 transliterate lower-case characters to upper-case
 — toupper, 853
 transliterate lower-case characters to upper-case
 — _toupper, 852
 transliterate lower-case wide-character code to
 upper-case — toupper, 856
 transliterate upper-case characters to lower-case
 — tolower, 851
 transliterate upper-case characters to lower-case
 — _tolower, 850
 transliterate upper-case wide-character code to
 lower-case — tolower, 855
 truncate — set a file to a specified length, 857
 tsearch — manage binary search trees, 860
 ttyname — find pathname of a terminal, 863
 ttyname_r — find pathname of a terminal, 863
 ttyslot — find the slot of the current user in the
 user accounting database, 865
 twalk — manage binary search trees, 860

U

ualarm — schedule signal after interval in
 microseconds, 866
 ucred_free — user credential functions, 867

ucred_get — user credential functions, 867
 ucred_getegid — user credential functions, 867
 ucred_geteuid — user credential functions, 867
 ucred_getgroups — user credential functions, 867
 ucred_getpflags — user credential functions, 867
 ucred_getpid — user credential functions, 867
 ucred_getprivset — user credential functions, 867
 ucred_getprojid — user credential functions, 867
 ucred_getrgid — user credential functions, 867
 ucred_getruid — user credential functions, 867
 ucred_getsgid — user credential functions, 867
 ucred_getsuid — user credential functions, 867
 ucred_getzoneid — user credential functions, 867
 ucred_size — user credential functions, 867
 ulltostr — string conversion routines, 750
 umem_alloc — fast, scalable memory allocation, 870
 umem_cache_alloc — allocation cache manipulation, 875
 umem_cache_create — allocation cache manipulation, 875
 umem_cache_destroy — allocation cache manipulation, 875
 umem_cache_free — allocation cache manipulation, 875
 umem_debug — debugging features of the umem library, 883
 umem_free — fast, scalable memory allocation, 870
 umem_nofail_callback — fast, scalable memory allocation, 870
 umem_zalloc — fast, scalable memory allocation, 870
 ungetc — push byte back into input stream, 886
 ungetwc — push wide-character code back into input stream, 887
 unlock a pseudo-terminal master/slave pair — unlockpt, 888
 unlock a read-write lock object — pthread_rwlock_unlock, 577
 unlock a spin lock object — pthread_spin_unlock, 594
 unlock address space, — munlockall, 386
 unlock memory pages, — munlock, 384
 unlockpt — unlock a pseudo-terminal master/slave pair, 888
 unordered — determine type of floating-point number, 323
 unsetenv — remove an environment variable, 889
 update stack boundaries — stack_setbounds, 708
 updwtmp — user accounting database functions, 285
 updwtmpx — user accounting database functions, 285
 user accounting database functions — endutent, 282
 user accounting database functions — endutxent, 285
 user accounting database functions — getutent, 282
 user accounting database functions — getutid, 282
 user accounting database functions — getutline, 282
 user accounting database functions — getutmp, 285
 user accounting database functions — getutmpx, 285
 user accounting database functions — getutxent, 285
 user accounting database functions — getutxid, 285
 user accounting database functions — getutxline, 285
 user accounting database functions — pututline, 282
 user accounting database functions — pututxline, 285
 user accounting database functions — setutent, 282
 user accounting database functions — setutxent, 285
 user accounting database functions — updwtmp, 285
 user accounting database functions — updwtmpx, 285
 user accounting database functions — utmpname, 282

- user accounting database functions —
 - utmpxname, 285
- user credential functions — `ucred_free`, 867
- user credential functions — `ucred_get`, 867
- user credential functions — `ucred_getegid`, 867
- user credential functions — `ucred_geteuid`, 867
- user credential functions —
 - `ucred_getgroups`, 867
- user credential functions —
 - `ucred_getpflags`, 867
- user credential functions — `ucred_getpid`, 867
- user credential functions —
 - `ucred_getprivset`, 867
- user credential functions —
 - `ucred_getprojid`, 867
- user credential functions — `ucred_getrgid`, 867
- user credential functions — `ucred_getruid`, 867
- user credential functions — `ucred_getsgid`, 867
- user credential functions — `ucred_getsuid`, 867
- user credential functions —
 - `ucred_getzoneid`, 867
- user credential functions — `ucred_size`, 867
- user IDs, get character-string representation —
 - `cuserid`, 79
- `usleep` — suspend execution for interval in microseconds, 890
- `utmpname` — user accounting database functions, 282
- `utmpx` file, find the slot of current user —
 - `ttyslot`, 865
- `utmpxname` — user accounting database functions, 285

V

- `vfprintf` — formatted output conversion, 471
- `vfscanf` — convert formatted input, 647
- `vfstab` file, — `getvfsent`, 289
- `vwprintf` — wide-character formatted output of a `stdarg` argument list, 891
- `vwscanf` — convert formatted wide-character input, 200
- virtual memory, optimizing usage of user mapped memory — `madvise`, 348
- visit registered `rctls` on current system —
 - `rctl_walk`, 618

- `vlfmt` — display error message in standard format and pass to logging and monitoring services, 892
- `vpfmt` — display error message in standard format and pass to logging and monitoring services, 894
- `vprintf` — formatted output conversion, 471
- `vscanf` — convert formatted input, 647
- `vsprintf` — formatted output conversion, 471
- `vsscanf` — convert formatted input, 647
- `vswprintf` — wide-character formatted output of a `stdarg` argument list, 891
- `vswscanf` — convert formatted wide-character input, 200
- `vsyslog` — log message with a `stdarg` argument list, 898
- `vwprintf` — wide-character formatted output of a `stdarg` argument list, 891
- `vwscanf` — convert formatted wide-character input, 200

W

- `wait` — wait for child process to stop or terminate, 902
- `wait` — wait for process to terminate or stop, 904
- `wait` on a condition —
 - `pthread_cond_reltimedwait_np`, 521
- `wait` on a condition —
 - `pthread_cond_timedwait`, 521
- `wait` on a condition — `pthread_cond_wait`, 521
- `waitpid` — wait for process to terminate or stop, 904
- `wait` for process to terminate or stop
 - `wait3`, 899
 - `wait4`, 899
- `wait` for thread termination —
 - `pthread_join`, 537
- `wait` for thread termination — `thr_join`, 829
- `wait` for transmission of output — `tcdrain`, 774
- `wait3` — wait for process to terminate or stop, 899, 904
- `wait4` — wait for process to terminate or stop, 899, 904
- `waitpid` — wait for child process to change state, 908

walk stack pointed to by ucontext —
 printstack, 910

walk stack pointed to by ucontext —
 walkcontext, 910

walkcontext — walk stack pointed to by
 ucontext, 910

watof — convert wide character string to
 floating-point number, 922

watoi — convert wide character string to long
 integer, 926

watol — convert wide character string to long
 integer, 926

watoll — convert wide character string to long
 integer, 926

wchar_t string, number conversion —
 wscanf, 955

wcrtomb — convert a wide-character code to a
 character (restartable), 915

wcscat — wide-character string operations, 933

wcschr — wide-character string operations, 934

wcscmp — wide-character string
 operations, 933

wcscoll — wide character string comparison
 using collating information, 917

wcscpy — wide-character string
 operations, 934

wcscspn — wide-character string
 operations, 935

wcsetno — get information on EUC
 codesets, 69

wcsftime — convert date and time to wide
 character string, 918

wcslen — wide-character string operations, 934

wcsncat — wide-character string
 operations, 933

wcsncmp — wide-character string
 operations, 934

wcsncpy — wide-character string
 operations, 934

wcspbrk — wide-character string
 operations, 935

wcsrchr — wide-character string
 operations, 934

wcrtombs — convert a wide-character string to
 a character string (restartable), 919

wcsspn — wide-character string
 operations, 935

wcsstr — find a wide-character substring, 921

wctod — convert wide character string to
 floating-point number, 922

wctof — convert wide character string to
 floating-point number, 922

wctoimax — convert wide-character string to
 integer type, 925

wcstok — wide-character string operations, 935

wcstol — convert wide character string to long
 integer, 926

wctold — convert wide character string to
 floating-point number, 922

wctoll — convert wide character string to long
 integer, 926

wctombs — convert a wide-character string to
 a character string, 929

wctoul — convert wide character string to
 unsigned long, 930

wcstoumax — convert wide-character string to
 integer type, 925

wcstring — wide-character string
 operations, 933

wcswcs — wide-character string
 operations, 935

wcswidth — number of column positions of a
 wide-character string, 937

wcsxfrm — wide character string
 transformation, 938

wctob — wide-character to single-byte
 conversion, 940

wctomb — convert a wide-character code to a
 character, 941

wctrans — define character mapping, 942

wctype — define character class, 943

wcwidth — number of column positions of a
 wide-character code, 944

wide-character formatted output of a stdarg
 argument list

- vfwprintf, 891
- vswprintf, 891
- vwprintf, 891

wide-character mapping — towctrans, 854

wide character string comparison using
 collating information

- wcscoll, 917
- wscoll, 917

wide-character string operations

- wcscat, 933
- wcschr, 934

