

# Bull

## AIX 5L Technical Reference Base Operating System and Extensions

Volume 2/2

AIX





# Bull

## AIX 5L Technical Reference Base Operating System and Extensions

Volume 2/2

AIX

---

Software

May 2003

**BULL CEDOC  
357 AVENUE PATTON  
B.P.20845  
49008 ANGERS CEDEX 01  
FRANCE**

**ORDER REFERENCE  
86 A2 48EF 02**

The following copyright notice protects this book under the Copyright laws of the United States of America and other countries which prohibit such actions as, but not limited to, copying, distributing, modifying, and making derivative works.

Copyright © Bull S.A. 1992, 2003

Printed in France

Suggestions and criticisms concerning the form, content, and presentation of this book are invited. A form is provided at the end of this book for this purpose.

To order additional copies of this book or other Bull Technical Publications, you are invited to use the Ordering Form also provided at the end of this book.

### **Trademarks and Acknowledgements**

We acknowledge the right of proprietors of trademarks mentioned in this book.

AIX<sup>®</sup> is a registered trademark of International Business Machines Corporation, and is being used under licence.

UNIX is a registered trademark in the United States of America and other countries licensed exclusively through the Open Group.

Linux is a registered trademark of Linus Torvalds.

---

# Contents

<b>About This Book</b> . . . . .	xiii
Who Should Use This Book . . . . .	xiii
Highlighting . . . . .	xiii
Case-Sensitivity in AIX . . . . .	xiii
ISO 9000 . . . . .	xiv
32-Bit and 64-Bit Support for the UNIX98 Specification . . . . .	xiv
Related Publications . . . . .	xiv
<b>Chapter 1. Base Operating System (BOS) Runtime Services (Q-Z)</b> . . . . .	1
qsort Subroutine . . . . .	1
quotactl Subroutine . . . . .	2
raise Subroutine . . . . .	4
rand or srand Subroutine . . . . .	5
rand_r Subroutine . . . . .	6
random, srandom, initstate, or setstate Subroutine . . . . .	7
ra_attachrset Subroutine . . . . .	9
ra_detachrset Subroutine . . . . .	11
ra_exec Subroutine . . . . .	12
ra_fork Subroutine . . . . .	14
ra_getrset Subroutine . . . . .	15
read, readx, readv, readvx, or pread Subroutine . . . . .	16
readdir_r Subroutine . . . . .	21
readlink Subroutine . . . . .	22
read_real_time or time_base_to_time Subroutine . . . . .	23
realpath Subroutine . . . . .	25
reboot Subroutine . . . . .	26
re_comp or re_exec Subroutine . . . . .	28
regcmp or regex Subroutine . . . . .	29
regcomp Subroutine . . . . .	32
regerror Subroutine . . . . .	33
regexec Subroutine . . . . .	35
regfree Subroutine . . . . .	38
retimerid Subroutine . . . . .	39
remainder, remainderf, or remainderl Subroutine . . . . .	39
remove Subroutine . . . . .	40
remquo, remquof, or remquol Subroutine . . . . .	41
rename Subroutine . . . . .	42
reset_malloc_log Subroutine . . . . .	44
revoke Subroutine . . . . .	45
rintf, rintl, or rint Subroutine . . . . .	46
round, roundf, or roundl Subroutine . . . . .	47
rmdir Subroutine . . . . .	48
rpmatch Subroutine . . . . .	49
RSiAddSetHot Subroutine . . . . .	50
RSiChangeFeed Subroutine . . . . .	53
RSiChangeHotFeed Subroutine . . . . .	54
RSiClose Subroutine . . . . .	55
RSiCreateHotSet Subroutine . . . . .	56
RSiCreateStatSet Subroutine . . . . .	57
RSiDelSetHot Subroutine . . . . .	58
RSiDelSetStat Subroutine . . . . .	59
RSiFirstCx Subroutine . . . . .	60
RSiFirstStat Subroutine . . . . .	62

RSiGetHotItem Subroutine . . . . .	63
RSiGetRawValue Subroutine . . . . .	65
RSiGetValue Subroutine . . . . .	66
RSiInit Subroutine . . . . .	68
RSiInstantiate Subroutine . . . . .	69
RSiInvite Subroutine . . . . .	70
RSiMainLoop Subroutine. . . . .	72
RSiNextCx Subroutine. . . . .	73
RSiNextStat Subroutine . . . . .	74
RSiOpen Subroutine . . . . .	75
RSiPathAddSetStat Subroutine . . . . .	78
RSiPathGetCx Subroutine . . . . .	79
RSiStartFeed Subroutine. . . . .	80
RSiStartHotFeed Subroutine . . . . .	81
RSiStatGetPath Subroutine . . . . .	82
RSiStopFeed Subroutine. . . . .	84
RSiStopHotFeed Subroutine . . . . .	85
rs_alloc Subroutine . . . . .	86
rs_discardname Subroutine . . . . .	87
rs_free Subroutine . . . . .	88
rs_getassociativity Subroutine . . . . .	89
rs_getinfo Subroutine . . . . .	90
rs_getnameattr Subroutine . . . . .	91
rs_getnamedrset Subroutine . . . . .	92
rs_getpartition Subroutine . . . . .	93
rs_getrad Subroutine . . . . .	94
rs_init Subroutine . . . . .	96
rs_numrads Subroutine . . . . .	96
rs_op Subroutine . . . . .	97
rs_registername Subroutine. . . . .	99
rs_setnameattr Subroutine. . . . .	101
rs_setpartition Subroutine . . . . .	103
rsqrt Subroutine . . . . .	104
rstat Subroutines . . . . .	105
scalbln, scalblnf, scalblnl, scalbn, scalbnf, scalbnl, or scalb Subroutine . . . . .	106
scandir or alphasort Subroutine . . . . .	107
scanf, fscanf, sscanf, or wscanf Subroutine . . . . .	109
sched_yield Subroutine . . . . .	115
select Subroutine . . . . .	115
semctl Subroutine . . . . .	120
semget Subroutine . . . . .	123
semop Subroutine. . . . .	125
setacldb or endacldb Subroutine . . . . .	127
setauthdb Subroutine . . . . .	128
setbuf, setvbuf, setbuffer, or setlinebuf Subroutine . . . . .	129
setcsmap Subroutine. . . . .	131
setgid, setrgid, setegid, setregid, or setgidx Subroutine . . . . .	132
setgroups Subroutine . . . . .	134
setjmp or longjmp Subroutine . . . . .	135
setlocale Subroutine . . . . .	136
setpcred Subroutine . . . . .	139
setpenv Subroutine . . . . .	142
setpgid or setpgrp Subroutine . . . . .	146
setpri Subroutine . . . . .	147
setpwdb or endpwdb Subroutine . . . . .	148
setroledb or endroledb Subroutine . . . . .	149

setuid Subroutine . . . . .	150
setuid, setruid, seteuid, setreuid or setuidx Subroutine . . . . .	151
setuserdb or enduserdb Subroutine . . . . .	153
sgetl or sputl Subroutine . . . . .	154
shmat Subroutine . . . . .	155
shmctl Subroutine . . . . .	158
shmdt Subroutine . . . . .	160
shmget Subroutine . . . . .	161
sigaction, sigvec, or signal Subroutine . . . . .	164
sigaltstack Subroutine . . . . .	174
sigemptyset, sigfillset, sigaddset, sigdelset, or sigismember Subroutine . . . . .	175
siginterrupt Subroutine . . . . .	177
sigbit Macro . . . . .	178
sigpending Subroutine . . . . .	179
sigprocmask, sigsetmask, or sigblock Subroutine . . . . .	179
sigqueue Subroutine . . . . .	181
sigset, sighold, sigelse, or sigignore Subroutine. . . . .	183
sigsetjmp or siglongjmp Subroutine . . . . .	186
sigstack Subroutine . . . . .	187
sigsuspend or sigpause Subroutine . . . . .	188
sigthreadmask Subroutine . . . . .	189
sigtimedwait and sigwaitinfo Subroutine . . . . .	191
sigwait Subroutine. . . . .	192
sin, sinf, or sinl Subroutine . . . . .	193
sinh, sinhf, or sinhl Subroutine . . . . .	194
sleep, nsleep or usleep Subroutine . . . . .	196
socketmark Subroutine . . . . .	197
SpmiAddSetHot Subroutine . . . . .	198
SpmiCreateHotSet . . . . .	201
SpmiCreateStatSet Subroutine . . . . .	202
SpmiDdsAddCx Subroutine . . . . .	203
SpmiDdsDelCx Subroutine . . . . .	204
SpmiDdsInit Subroutine . . . . .	205
SpmiDelSetHot Subroutine . . . . .	207
SpmiDelSetStat Subroutine . . . . .	208
SpmiExit Subroutine . . . . .	210
SpmiFirstCx Subroutine. . . . .	210
SpmiFirstHot Subroutine . . . . .	211
SpmiFirstStat Subroutine . . . . .	213
SpmiFirstVals Subroutine . . . . .	214
SpmiFreeHotSet Subroutine . . . . .	215
SpmiFreeStatSet Subroutine . . . . .	216
SpmiGetCx Subroutine . . . . .	217
SpmiGetHotSet Subroutine . . . . .	218
SpmiGetStat Subroutine . . . . .	219
SpmiGetStatSet Subroutine . . . . .	221
SpmiGetValue Subroutine . . . . .	222
Spmilnit Subroutine . . . . .	223
SpmiInstantiate Subroutine . . . . .	225
SpmiNextCx Subroutine . . . . .	226
SpmiNextHot Subroutine . . . . .	227
SpmiNextHotItem Subroutine. . . . .	228
SpmiNextStat Subroutine . . . . .	230
SpmiNextVals Subroutine . . . . .	232
SpmiNextValue Subroutine . . . . .	232
SpmiPathAddSetStat Subroutine . . . . .	234

SpmiPathGetCx Subroutine . . . . .	236
SpmiStatGetPath Subroutine . . . . .	237
sqrt, sqrtf, or sqrtl Subroutine . . . . .	238
src_err_msg Subroutine . . . . .	240
src_err_msg_r Subroutine . . . . .	241
srcrrqs Subroutine . . . . .	241
srcrrqs_r Subroutine . . . . .	243
srcsbuf Subroutine . . . . .	244
srcsbuf_r Subroutine . . . . .	247
srcsrpy Subroutine . . . . .	250
srcsrqt Subroutine . . . . .	253
srcsrqt_r Subroutine . . . . .	256
srcstat Subroutine . . . . .	259
srcstat_r Subroutine . . . . .	262
srcstathdr Subroutine . . . . .	264
srcstattxt Subroutine . . . . .	265
srcstattxt_r Subroutine . . . . .	265
srcstop Subroutine . . . . .	266
srcstrt Subroutine . . . . .	269
ssignal or gsignal Subroutine . . . . .	271
statacl or fstatacl Subroutine . . . . .	272
statfs, fstatfs, or ustat Subroutine . . . . .	274
statvfs or fstatvfs Subroutine . . . . .	276
statx, stat, lstat, fstatx, fstat, fullstat, ffullstat, stat64, lstat64, or fstat64 Subroutine . . . . .	277
strcat, strncat, strxfrm, strcpy, strncpy, or strdup Subroutine . . . . .	281
strcmp, strncmp, strcasecmp, strncasecmp, or strcoll Subroutine . . . . .	283
strerror Subroutine . . . . .	285
strfmon Subroutine . . . . .	286
strftime Subroutine . . . . .	288
strlen, strchr, strrchr, strpbrk, strspn, strcspn, strstr, strtok, or strsep Subroutine . . . . .	291
strncollen Subroutine . . . . .	294
strtof, strtod, or strtold Subroutine . . . . .	295
strtoimax or strtoumax Subroutine . . . . .	297
strtok_r Subroutine . . . . .	298
strtol, strtoul, strtoll, strtoull, or atoi Subroutine . . . . .	299
strptime Subroutine . . . . .	301
stty or gtty Subroutine . . . . .	303
swab Subroutine . . . . .	304
swapoff Subroutine . . . . .	305
swapon Subroutine . . . . .	306
swapqry Subroutine . . . . .	307
symlink Subroutine . . . . .	308
sync Subroutine . . . . .	310
_sync_cache_range Subroutine . . . . .	311
sysconf Subroutine . . . . .	311
sysconfig Subroutine . . . . .	315
SYS_CFGDD sysconfig Operation . . . . .	316
SYS_CFGKMOD sysconfig Operation . . . . .	318
SYS_GETLPAR_INFO sysconfig Operation . . . . .	319
SYS_GETPARMS sysconfig Operation . . . . .	320
SYS_KLOAD sysconfig Operation . . . . .	321
SYS_KULOAD sysconfig Operation . . . . .	323
SYS_QDVSX sysconfig Operation . . . . .	325
SYS_QUERYLOAD sysconfig Operation . . . . .	326
SYS_SETPARMS sysconfig Operation . . . . .	327
SYS_SINGLELOAD sysconfig Operation . . . . .	328

syslog, openlog, closelog, or setlogmask Subroutine . . . . .	329
syslog_r, openlog_r, closelog_r, or setlogmask_r Subroutine . . . . .	332
sys_parm Subroutine. . . . .	336
system Subroutine . . . . .	337
tan, tanf, or tanl Subroutine . . . . .	339
tanh, tanhf, or tanhl Subroutine . . . . .	340
tcb Subroutine . . . . .	341
tcdrain Subroutine. . . . .	342
tcfLOW Subroutine . . . . .	343
tcfLUSH Subroutine . . . . .	344
togetAttr Subroutine . . . . .	346
togetpgrp Subroutine . . . . .	347
tcSENDBREAK Subroutine. . . . .	348
tcSETATTR Subroutine . . . . .	349
tcSETPGRP Subroutine . . . . .	351
termdef Subroutine . . . . .	352
test_and_set Subroutine . . . . .	353
tgamma, tgammaf, or tgammaL Subroutine . . . . .	354
times Subroutine . . . . .	355
timezone Subroutine . . . . .	357
thread_post Subroutine . . . . .	358
thread_post_many Subroutine . . . . .	359
thread_self Subroutine . . . . .	360
thread_setsched Subroutine . . . . .	360
thread_wait Subroutine . . . . .	362
tmpfile Subroutine . . . . .	363
tmpnam or tmpnam Subroutine . . . . .	363
towctrans Subroutine. . . . .	365
towlower Subroutine . . . . .	366
towupper Subroutine . . . . .	367
t_rcvreldata Subroutine . . . . .	367
t_rcvv Subroutine . . . . .	369
t_rcvvudata Subroutine . . . . .	371
t_sndv Subroutine . . . . .	373
t_sndreldata Subroutine. . . . .	376
t_sndvudata Subroutine. . . . .	377
t_sysconf Subroutine. . . . .	379
trc_close Subroutine . . . . .	380
trc_find_first, trc_find_next, and trc_compare Subroutine . . . . .	381
trc_free Subroutine . . . . .	385
trc_hkemptysset, trc_hkfillset, trc_hkaddset, trc_hkdelset, and trc_hkisset Subroutine . . . . .	386
trc_hookname Subroutine . . . . .	387
trc_libcntl Subroutine. . . . .	388
trc_loginfo Subroutine . . . . .	390
trc_open Subroutine . . . . .	392
trc_perror Subroutine. . . . .	394
trc_read Subroutine . . . . .	395
trc_seek and trc_tell Subroutine. . . . .	398
trc_strerror Subroutine . . . . .	399
trcgen or trcgent Subroutine . . . . .	400
trchook, utrchook, trchook64, and utrhook64 Subroutine . . . . .	402
trcoff Subroutine . . . . .	404
trcon Subroutine . . . . .	404
trcstart Subroutine. . . . .	405
trcstop Subroutine. . . . .	406
trunc, truncf, or truncL Subroutine . . . . .	406

truncate, truncate64, ftruncate, or ftruncate64 Subroutine . . . . .	407
tsearch, tdelete, tfind or twalk Subroutine . . . . .	410
ttylock, ttywait, ttyunlock, or ttylocked Subroutine . . . . .	412
ttyname or isatty Subroutine . . . . .	413
ttyslot Subroutine . . . . .	415
ulimit Subroutine . . . . .	415
umask Subroutine . . . . .	418
umount or uvmount Subroutine . . . . .	419
uname or unamex Subroutine . . . . .	420
ungetc or ungetwc Subroutine . . . . .	422
unlink Subroutine . . . . .	423
unload Subroutine . . . . .	424
unlockpt Subroutine . . . . .	426
usrinfo Subroutine . . . . .	426
utimes or utime Subroutine . . . . .	428
varargs Macros . . . . .	429
vfscanf, vscanf, or vsscanf Subroutine . . . . .	432
vwscanf, vswscanf, or vwscanf Subroutine . . . . .	433
vwprintf, vwprintf Subroutine . . . . .	433
vmgetinfo Subroutine . . . . .	434
vmount or mount Subroutine . . . . .	435
vsprintf Subroutine . . . . .	439
vwsprintf Subroutine . . . . .	439
wait, waitpid, wait3, or wait364 Subroutine . . . . .	440
waitid Subroutine . . . . .	443
wcscat, wcschr, wcsncmp, wcsncpy, or wcsncpy Subroutine . . . . .	444
wscoll Subroutine . . . . .	446
wcsftime Subroutine . . . . .	447
wcsid Subroutine . . . . .	448
wcslen Subroutine . . . . .	449
wcsncat, wcsncmp, or wcsncpy Subroutine . . . . .	449
wcspbrk Subroutine . . . . .	451
wcsrchr Subroutine . . . . .	451
wcsrtombs Subroutine . . . . .	452
wcsspn Subroutine . . . . .	453
wcsstr Subroutine . . . . .	454
wcstod, wcstof, or wcstold Subroutine . . . . .	455
wcstoimax or wcstoumax Subroutine . . . . .	457
wcstok Subroutine . . . . .	458
wcstol or wcstoll Subroutine . . . . .	459
wcstombs Subroutine . . . . .	461
wcstoul or wcstoull Subroutine . . . . .	463
wcswcs Subroutine . . . . .	464
wcswidth Subroutine . . . . .	465
wcsxfrm Subroutine . . . . .	466
wctob Subroutine . . . . .	468
wctomb Subroutine . . . . .	468
wctrans Subroutine . . . . .	469
wctype or get_wctype Subroutine . . . . .	470
wcwidth Subroutine . . . . .	471
wlm_assign Subroutine . . . . .	473
wlm_change_class Subroutine . . . . .	475
wlm_check subroutine . . . . .	476
wlm_classify Subroutine . . . . .	477
wlm_class2key Subroutine . . . . .	479
wlm_create_class Subroutine . . . . .	480

wlm_delete_class Subroutine . . . . .	482
wlm_endkey Subroutine . . . . .	483
wlm_get_bio_stats subroutine . . . . .	484
wlm_get_info Subroutine . . . . .	486
wlm_get_procinfo Subroutine . . . . .	488
wlm_init_class_definition Subroutine . . . . .	489
wlm_initialize Subroutine . . . . .	490
wlm_initkey Subroutine . . . . .	491
wlm_key2class Subroutine . . . . .	492
wlm_load Subroutine . . . . .	493
wlm_read_classes Subroutine . . . . .	495
wlm_set Subroutine . . . . .	497
wlm_set_tag Subroutine . . . . .	499
wmemchr Subroutine . . . . .	500
wmemcmp Subroutine . . . . .	501
wmemcpy Subroutine . . . . .	501
wmemmove Subroutine . . . . .	502
wmemset Subroutine . . . . .	503
wordexp Subroutine . . . . .	503
wordfree Subroutine . . . . .	506
write, writex, writev, writevx or pwrite Subroutines . . . . .	506
wstring Subroutine . . . . .	511
wstrtod or watof Subroutine . . . . .	514
wstrtol, watol, or watoi Subroutine . . . . .	515
xcrypt_key_setup, xcrypt_encrypt, xcrypt_decrypt, xcrypt_hash, xcrypt_malloc, xcrypt_free, xcrypt_printb, xcrypt_btoa and xcrypt_randbuff Subroutine . . . . .	516
yield Subroutine . . . . .	520
<b>Chapter 2. Curses Subroutines . . . . .</b>	<b>523</b>
addch, mvaddch, mvwaddch, or waddch Subroutine . . . . .	523
addnstr, addstr, mvaddnstr, mvaddstr, mvwaddnstr, mvwaddstr, waddnstr, or waddstr Subroutine . . . . .	524
attroff, attron, attrset, wattroff, attron, or wattrset Subroutine . . . . .	526
attron or wattron Subroutine . . . . .	528
attrset or wattrset Subroutine . . . . .	528
baudrate Subroutine . . . . .	529
beep Subroutine . . . . .	530
box Subroutine . . . . .	531
can_change_color, color_content, has_colors,init_color, init_pair, start_color or pair_content Subroutine . . . . .	532
cbreak, nocbreak, noraw, or raw Subroutine . . . . .	535
clear, erase, wclear or werase Subroutine . . . . .	536
clearok, idlok, leaveok, scrollok, setscreg or wsetscreg Subroutine . . . . .	537
clrtoebot or wclrtoebot Subroutine . . . . .	540
clrtoeol or wclrtoeol Subroutine . . . . .	541
color_content Subroutine . . . . .	542
copywin Subroutine . . . . .	543
curs_set Subroutine . . . . .	544
def_prog_mode, def_shell_mode, reset_prog_mode or reset_shell_mode Subroutine . . . . .	545
def_shell_mode Subroutine . . . . .	546
del_curterm, restartterm, set_curterm, or setupterm Subroutine . . . . .	547
delay_output Subroutine . . . . .	549
delch, mvdelch, mvwdelch or wdelch Subroutine . . . . .	550
deleteln or wdeleteln Subroutine . . . . .	551
delwin Subroutine . . . . .	552
echo or noecho Subroutine . . . . .	552
echochar or wechochar Subroutines . . . . .	553

endwin Subroutine . . . . .	554
erase or werase Subroutine . . . . .	555
erasechar, erasewchar, killchar, and killwchar Subroutine . . . . .	556
filter Subroutine. . . . .	557
flash Subroutine . . . . .	557
flushinp Subroutine . . . . .	558
garbagedlines Subroutine . . . . .	559
getbegyx, getmaxyx, getparyx, or getyx Subroutine . . . . .	560
getch, mvgetch, mvwgetch, or wgetch Subroutine . . . . .	561
getmaxyx Subroutine. . . . .	565
getnstr, getstr, mvgetnstr, mvgetstr, mvwgetnstr, mvwgetstr, wgetnstr, or wgetstr Subroutine . . . . .	566
getsyx Subroutine . . . . .	568
getyx Macro . . . . .	569
halfdelay Subroutine . . . . .	570
has_colors Subroutine . . . . .	570
has_ic and has_il Subroutine. . . . .	571
has_il Subroutine . . . . .	572
idlok Subroutine . . . . .	573
inch, mvinch, mvwinch, or winch Subroutine . . . . .	574
init_color Subroutine . . . . .	575
init_pair Subroutine . . . . .	576
initscr and newterm Subroutine . . . . .	577
insch, mvinsch, mvwinsch, or winsch Subroutine . . . . .	578
insertln or winsertln Subroutine . . . . .	579
intrflush Subroutine . . . . .	580
keyname, key_name Subroutine . . . . .	581
keypad Subroutine . . . . .	582
killchar or killwchar Subroutine . . . . .	583
_lazySetErrorHandler Subroutine . . . . .	584
leaveok Subroutine . . . . .	585
longname Subroutine . . . . .	586
makenew Subroutine. . . . .	587
meta Subroutine . . . . .	588
move or wmove Subroutine . . . . .	589
mvcur Subroutine . . . . .	589
mvwin Subroutine . . . . .	591
newpad, pnoutrefresh, prefresh, or subpad Subroutine . . . . .	592
newterm Subroutine . . . . .	594
derwin, newwin, or subwin Subroutine . . . . .	596
nl or nonl Subroutine. . . . .	598
nodelay Subroutine . . . . .	599
notimeout, timeout, wtimeout Subroutine . . . . .	599
overlay or overwrite Subroutine . . . . .	601
pair_content Subroutine. . . . .	602
prefresh or pnoutrefresh Subroutine . . . . .	603
printw, wprintw, mvprintw, or mvwprintw Subroutine . . . . .	604
putp, tputs Subroutine . . . . .	606
raw or noraw Subroutine . . . . .	607
refresh or wrefresh Subroutine . . . . .	608
reset_prog_mode Subroutine. . . . .	609
reset_shell_mode Subroutine. . . . .	610
resetterm Subroutine. . . . .	610
resetty, savetty Subroutine. . . . .	611
restartterm Subroutine . . . . .	612
ripline Subroutine . . . . .	612
savetty Subroutine . . . . .	614

scanw, wscanw, mvscanw, or mvwscanw Subroutine . . . . .	614
scr_dump, scr_init, scr_restore, scr_set Subroutine . . . . .	615
scr_init Subroutine . . . . .	617
scr_restore Subroutine . . . . .	618
screl, scroll, wscrl Subroutine . . . . .	619
scrollok Subroutine . . . . .	620
set_curterm Subroutine . . . . .	621
setscrreg or wsetscrreg Subroutine . . . . .	622
setsyx Subroutine . . . . .	623
set_term Subroutine . . . . .	624
setupterm Subroutine . . . . .	624
_showstring Subroutine . . . . .	626
slk_atroff, slk_attr_off, slk_attron, slk_attrset, slk_attr_set, slk_clear, slk_color, slk_init, slk_label, slk_noutrefresh, slk_refresh, slk_restore, slk_set, slk_touch, slk_wset, Subroutine . . . . .	627
slk_init Subroutine . . . . .	630
slk_label Subroutine . . . . .	631
slk_noutrefresh Subroutine . . . . .	631
slk_refresh Subroutine . . . . .	632
slk_restore Subroutine . . . . .	633
slk_set Subroutine . . . . .	633
slk_touch Subroutine . . . . .	634
standend, standout, wstandend, or wstandout Subroutine . . . . .	635
start_color Subroutine . . . . .	636
subpad Subroutine . . . . .	637
subwin Subroutine . . . . .	638
tgetent, tgetflag, tgetnum, tgetstr, or tgoto Subroutine . . . . .	640
tgetflag Subroutine . . . . .	641
tgetnum Subroutine . . . . .	642
tgetstr Subroutine . . . . .	642
tgoto Subroutine . . . . .	643
tigetflag, tigetnum, tigetstr, or tparm Subroutine . . . . .	644
tigetnum Subroutine . . . . .	646
tigetstr Routine . . . . .	647
is_linetouched, is_wintouched, touchline, touchwin, untouchwin, or wtouchin Subroutine . . . . .	648
touchoverlap Subroutine . . . . .	649
touchwin Subroutine . . . . .	650
tparm Subroutine . . . . .	651
tputs Subroutine . . . . .	652
typeahead Subroutine . . . . .	653
unctrl Subroutine . . . . .	654
ungetch, unget_wch Subroutine . . . . .	655
vidattr, vid_attr, vidputs, or vid_puts Subroutine . . . . .	656
doupdate, refresh, wnoutrefresh, or wrefresh Subroutines . . . . .	657
<b>Chapter 3. FORTRAN Basic Linear Algebra Subroutines (BLAS) . . . . .</b>	<b>661</b>
SDOT or DDOT Function . . . . .	661
CDOTC or ZDOTC Function . . . . .	661
CDOTU or ZDOTU Function . . . . .	662
SAXPY, DAXPY, CAXPY, or ZAXPY Subroutine . . . . .	663
SROTG, DROTG, CROTG, or ZROTG Subroutine . . . . .	663
SROT, DROT, CSROT, or ZDROT Subroutine . . . . .	664
SCOPY, DCOPY, CCOPY, or ZCOPY Subroutine . . . . .	665
SSWAP, DSWAP, CSWAP, or ZSWAP Subroutine . . . . .	666
SNRM2, DNRM2, SCNRM2, or DZNRM2 Function . . . . .	667
SASUM, DASUM, SCASUM, or DZASUM Function . . . . .	668
SSCAL, DSCAL, CSSCAL, CSCAL, ZDSCAL, or ZSCAL Subroutine . . . . .	668

ISAMAX, IDAMAX, ICAMAX, or IZAMAX Function . . . . .	669
SDSDOT Function . . . . .	670
SROTM or DROTM Subroutine . . . . .	671
SROTMG or DROTMG Subroutine . . . . .	672
SGEMV, DGEMV, CGEMV, or ZGEMV Subroutine . . . . .	673
SGBMV, DGBMV, CGBMV, or ZGBMV Subroutine . . . . .	674
CHEMV or ZHEMV Subroutine . . . . .	676
CHBMV or ZHBMV Subroutine . . . . .	677
CHPMV or ZHPMV Subroutine . . . . .	679
SSYMV or DSYMV Subroutine . . . . .	680
SSBMV or DSBMV Subroutine . . . . .	681
SSPMV or DSPMV Subroutine . . . . .	682
STRMV, DTRMV, CTRMV, or ZTRMV Subroutine . . . . .	683
STBMV, DTBMV, CTBMV, or ZTBMV Subroutine . . . . .	685
STPMV, DTPMV, CTPMV, or ZTPMV Subroutine . . . . .	687
STRSV, DTRSV, CTRSV, or ZTRSV Subroutine . . . . .	688
STBSV, DTBSV, CTBSV, or ZTBSV Subroutine . . . . .	690
STPSV, DTPSV, CTPSV, or ZTPSV Subroutine . . . . .	692
SGER or DGER Subroutine . . . . .	694
CGERU or ZGERU Subroutine . . . . .	695
CGERC or ZGERC Subroutine . . . . .	695
CHER or ZHER Subroutine . . . . .	696
CHPR or ZHPR Subroutine . . . . .	697
CHER2 or ZHER2 Subroutine . . . . .	698
CHPR2 or ZHPR2 Subroutine . . . . .	699
SSYR or DSYR Subroutine . . . . .	700
SSPR or DSPR Subroutine . . . . .	701
SSYR2 or DSYR2 Subroutine . . . . .	702
SSPR2 or DSPR2 Subroutine . . . . .	704
SGEMM, DGEMM, CGEMM, or ZGEMM Subroutine . . . . .	705
SSYMM, DSYMM, CSYMM, or ZSYMM Subroutine . . . . .	706
CHEMM or ZHEMM Subroutine . . . . .	708
SSYRK, DSYRK, CSYRK, or ZSYRK Subroutine . . . . .	710
CHERK or ZHERK Subroutine . . . . .	711
SSYR2K, DSYR2K, CSYR2K, or ZSYR2K Subroutine . . . . .	713
CHER2K or ZHER2K Subroutine . . . . .	714
STRMM, DTRMM, CTRMM, or ZTRMM Subroutine . . . . .	716
STRSM, DTRSM, CTRSM, or ZTRSM Subroutine . . . . .	718
<b>Appendix A. Base Operating System Error Codes for Services That Require Path-Name Resolution . . . . .</b>	<b>721</b>
<b>Appendix B. ODM Error Codes . . . . .</b>	<b>723</b>
<b>Appendix C. Notices . . . . .</b>	<b>725</b>
Trademarks . . . . .	726
<b>Index . . . . .</b>	<b>727</b>

---

## About This Book

This book provides information on application programming interfaces to the operating system.

This book is part of the six-volume technical reference set, *AIX 5L Version 5.2 Technical Reference*, that provides information on system calls, kernel extension calls, and subroutines in the following volumes:

- *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1* and *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 2* provide information on system calls, subroutines, functions, macros, and statements associated with base operating system runtime services.
- *AIX 5L Version 5.2 Technical Reference: Communications Volume 1* and *AIX 5L Version 5.2 Technical Reference: Communications Volume 2* provide information on entry points, functions, system calls, subroutines, and operations related to communications services.
- *AIX 5L Version 5.2 Technical Reference: Kernel and Subsystems Volume 1* and *AIX 5L Version 5.2 Technical Reference: Kernel and Subsystems Volume 2* provide information about kernel services, device driver operations, file system operations, subroutines, the configuration subsystem, the communications subsystem, the low function terminal (LFT) subsystem, the logical volume subsystem, the M-audio capture and playback adapter subsystem, the printer subsystem, the SCSI subsystem, and the serial DASD subsystem.

This edition supports the release of AIX 5L Version 5.2 with the 5200-01 Recommended Maintenance package. Any specific references to this maintenance package are indicated as *AIX 5.2 with 5200-01*.

---

## Who Should Use This Book

This book is intended for experienced C programmers. To use the book effectively, you should be familiar with commands, system calls, subroutines, file formats, and special files.

---

## Highlighting

The following highlighting conventions are used in this book:

<b>Bold</b>	Identifies commands, subroutines, keywords, files, structures, directories, and other items whose names are predefined by the system. Also identifies graphical objects such as buttons, labels, and icons that the user selects.
<i>Italics</i>	Identifies parameters whose actual names or values are to be supplied by the user.
Monospace	Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of portions of program code similar to what you might write as a programmer, messages from the system, or information you should actually type.

---

## Case-Sensitivity in AIX

Everything in the AIX operating system is case-sensitive, which means that it distinguishes between uppercase and lowercase letters. For example, you can use the **ls** command to list files. If you type **LS**, the system responds that the command is "not found." Likewise, **FILEA**, **FiLea**, and **filea** are three distinct file names, even if they reside in the same directory. To avoid causing undesirable actions to be performed, always ensure that you use the correct case.

---

## ISO 9000

ISO 9000 registered quality systems were used in the development and manufacturing of this product.

---

## 32-Bit and 64-Bit Support for the UNIX98 Specification

Beginning with Version 4.3, the operating system is designed to support The Open Group's UNIX98 Specification for portability of UNIX-based operating systems. Many new interfaces, and some current ones, have been added or enhanced to meet this specification, making Version 4.3 even more open and portable for applications.

At the same time, compatibility with previous releases of the operating system is preserved. This is accomplished by the creation of a new environment variable, which can be used to set the system environment on a per-system, per-user, or per-process basis.

To determine the proper way to develop a UNIX98-portable application, you may need to refer to The Open Group's UNIX98 Specification, which can be obtained on a CD-ROM by ordering *Go Solo 2: The Authorized Guide to Version 2 of the Single UNIX Specification*, a book which includes The Open Group's UNIX98 Specification on a CD-ROM.

---

## Related Publications

The following books contain information about or related to application programming interfaces:

- *AIX 5L Version 5.2 System Management Guide: Operating System and Devices*
- *AIX 5L Version 5.2 System Management Guide: Communications and Networks*
- *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*
- *AIX 5L Version 5.2 Communications Programming Concepts*
- *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*
- *AIX 5L Version 5.2 Files Reference*

---

# Chapter 1. Base Operating System (BOS) Runtime Services (Q-Z)

---

## qsort Subroutine

### Purpose

Sorts a table of data in place.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
void qsort (Base, NumberOfElements, Size, ComparisonPointer)
void * Base;
size_t NumberOfElements, Size;
int (*ComparisonPointer)(const void*, const void*);
```

### Description

The **qsort** subroutine sorts a table of data in place. It uses the quicker-sort algorithm.

### Parameters

<i>Base</i>	Points to the element at the base of the table.
<i>NumberOfElements</i>	Specifies the number of elements in the table.
<i>Size</i>	Specifies the size of each element.
<i>ComparisonPointer</i>	Points to the comparison function, which is passed two parameters that point to the objects being compared. The <b>qsort</b> subroutine sorts the array in ascending order according to the comparison function.

### Return Values

The comparison function compares its parameters and returns a value as follows:

- If the first parameter is less than the second parameter, the *ComparisonPointer* parameter returns a value less than 0.
- If the first parameter is equal to the second parameter, the *ComparisonPointer* parameter returns 0.
- If the first parameter is greater than the second parameter, the *ComparisonPointer* parameter returns a value greater than 0.

Because the comparison function need not compare every byte, the elements can contain arbitrary data in addition to the values being compared.

**Note:** If two items are the same when compared, their order in the output of this subroutine is unpredictable.

The pointer to the base of the table should be of type pointer-to-element, and cast to type pointer-to-character.

## Related Information

The **bsearch** subroutine, **lsearch** subroutine.

Searching and Sorting Example Program, Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## quotactl Subroutine

### Purpose

Manipulates disk quotas.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <jfs/quotactl.h>
```

```
int quotactl (Path, Cmd, ID, Addr)
int Cmd, ID;
char * Addr, * Path;
```

### Description

The **quotactl** subroutine enables, disables, and manipulates disk quotas for file systems on which quotas have been enabled.

Currently, disk quotas are supported only by the Journaled File System (JFS).

### Parameters

*Path* Specifies the path name of any file within the mounted file system to which the quota control command is to be applied.

*Cmd* Specifies the quota control command to be applied and whether it is applied to a user or group quota.

For JFSs, the *Cmd* parameter can be constructed through use of the **QCMD**(*Cmd*, *type*) macro contained within the **jfs/quota.h** file. The *Cmd* parameter specifies the quota control command. The *type* parameter specifies either user (**USRQUOTA**) or group (**GRPQUOTA**) quota type.

The valid JFS specific quota control values for the *Cmd* parameter are:

#### **Q\_QUOTAON**

Enables disk quotas for the file system specified by the *Path* parameter. The *Addr* parameter specifies a file from which to take the quotas. The quota file must exist; it is normally created with the **quotacheck** command. The *ID* parameter is unused. Root user authority is required to enable quotas.

#### **Q\_QUOTAOFF**

Disables disk quotas for the file system specified by the *Path* parameter. The *Addr* and *ID* arguments are unused. Root user authority is required to disable quotas.

#### **Q\_GETQUOTA**

Gets disk quota limits and current usage for a user or group specified by the *ID* parameter. The *Addr* parameter points to a **dqblk** buffer to hold the returned information. The **dqblk** structure is defined in the **jfs/quota.h** file. Root user authority is required if the *ID* value is not the current ID of the caller.

#### **Q\_SETQUOTA**

Sets disk quota limits for the user or group specified by the *ID* parameter. The *Addr* parameter points to a **dqblk** buffer containing the new quota limits. The **dqblk** structure is defined in the **jfs/quota.h** file. Root user authority is required to set quotas.

#### **Q\_SETUSE**

Sets disk usage limits for the user or group specified by the *ID* parameter. The *Addr* parameter points to a **dqblk** buffer containing the new usage limits. The **dqblk** structure is defined in the **jfs/quota.h** file. Root user authority is required to set disk usage limits.

*ID* Specifies the user or group ID to which the quota control command applies. The *ID* parameter is interpreted by the specified quota type. The JFS file system supports quotas for IDs within the range of **MINDQUID** through **MAXDQID**.

*Addr* Points to the address of an optional, command specific, data structure that is copied in or out of the system. The interpretation of the *Addr* parameter for each quota control command is given above.

## Return Values

A successful call returns 0, otherwise the value -1 is returned and the **errno** global variable indicates the reason for the failure.

## Error Codes

A **quotactl** subroutine will fail when one of the following occurs:

<b>EACCES</b>	In the <b>Q_QUOTAON</b> command, the quota file is not a regular file.
<b>EACCES</b>	Search permission is denied for a component of a path prefix.
<b>EFAULT</b>	An invalid <i>Addr</i> parameter is supplied; the associated structure could not be copied in or out of the kernel.
<b>EFAULT</b>	The <i>Path</i> parameter points outside the process's allocated address space.
<b>EINVAL</b>	The specified quota control command or quota type is invalid.
<b>EINVAL</b>	Path name contains a character with the high-order bit set.
<b>EINVAL</b>	The <i>ID</i> parameter is outside of the supported range ( <b>MINDQUID</b> through <b>MAXDQID</b> ).
<b>EIO</b>	An I/O error occurred while reading from or writing to a file containing quotas.
<b>ELOOP</b>	Too many symbolic links were encountered in translating a path name.
<b>ENAMETOOLONG</b>	A component of either path name exceeded 255 characters, or the entire length of either path name exceeded 1023 characters.

<b>ENOENT</b>	A file name does not exist.
<b>ENOTBLK</b>	Mounted file system is not a block device.
<b>ENOTDIR</b>	A component of a path prefix is not a directory.
<b>EOPNOTSUPP</b>	The file system does not support quotas.
<b>EPERM</b>	The quota control commands is privileged and the caller did not have root user authority.
<b>EROFS</b>	In the <b>Q_QUOTAON</b> command, the quota file resides on a read-only file system.
<b>EUSERS</b>	The in-core quota table cannot be expanded.

## Related Information

The **quotacheck** command.

Disk Quota System Overview in *AIX 5L Version 5.2 Security Guide*.

## raise Subroutine

### Purpose

Sends a signal to the currently running program.

### Libraries

Standard C Library (**libc.a**)

Threads Library (**libpthreads.a**)

### Syntax

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

```
int raise ( Signal )
int Signal;
```

### Description

The **raise** subroutine sends the signal specified by the *Signal* parameter to the executing process or thread, depending if the POSIX threads API (the **libpthreads.a** library) is used or not. When the program is not linked with the threads library, the **raise** subroutine sends the signal to the calling process as follows:

```
return kill(getpid(), Signal);
```

When the program is linked with the threads library, the **raise** subroutine sends the signal to the calling thread as follows:

```
return pthread_kill(pthread_self(), Signal);
```

When using the threads library, it is important to ensure that the threads library is linked before the standard C library.

### Parameter

*Signal*      Specifies a signal number.

## Return Values

Upon successful completion of the **raise** subroutine, a value of 0 is returned. Otherwise, a nonzero value is returned, and the **errno** global variable is set to indicate the error.

## Error Code

**EINVAL** The value of the **sig** argument is an invalid signal number

## Related Information

The **\_exit** subroutine, **kill** subroutine, **pthread\_kill** subroutine, **sigaction** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine.

Signal Management in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs* provides more information about signal management in multi-threaded processes.

---

## rand or srand Subroutine

### Purpose

Generates pseudo-random numbers.

### Library

Standard C Library (**libc.a**)

### Syntax

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

void srand ( Seed)
unsigned int Seed;
```

### Description

**Attention:** Do not use the **rand** subroutine in a multithreaded environment. See the multithread alternative in the **rand\_r** (“rand\_r Subroutine” on page 6) subroutine article.

The **rand** subroutine generates a pseudo-random number using a multiplicative congruential algorithm. The random-number generator has a period of  $2^{32}$ , and it returns successive pseudo-random numbers in the range from 0 through  $(2^{15}) - 1$ .

The **srand** subroutine resets the random-number generator to a new starting point. It uses the *Seed* parameter as a seed for a new sequence of pseudo-random numbers to be returned by subsequent calls to the **rand** subroutine. If you then call the **srand** subroutine with the same seed value, the **rand** subroutine repeats the sequence of pseudo-random numbers. When you call the **rand** subroutine before making any calls to the **srand** subroutine, it generates the same sequence of numbers that it would if you first called the **srand** subroutine with a seed value of 1.

**Note:** The **rand** subroutine is a simple random-number generator. Its spectral properties, a mathematical measurement of randomness, are somewhat limited. See the **drand48** subroutine or the **random** subroutine for more elaborate random-number generators that have greater spectral properties.

## Parameter

*Seed* Specifies an initial seed value.

## Return Values

Upon successful completion, the **rand** subroutine returns the next random number in sequence. The **srand** subroutine returns no value.

There are better random number generators, as noted above; however, the **rand** and **srand** subroutines are the interfaces defined for the ANSI C library.

## Example

The following functions define the semantics of the **rand** and **srand** subroutines, and are included here to facilitate porting applications from different implementations:

```
static unsigned int next = 1;
int rand( )
{
    next = next
    *
    1103515245 + 12345;
    return ((next >>16) & 32767);
}

void srand (Seed)
```

```
unsigned
int Seed;
{
    next = Seed;
}
```

## Related Information

The **drand48**, **erand48**, **lrand48**, **nrand48**, **mrand48**, **jrand48**, **srand48**, **seed48**, or **lcg48** subroutine, **random**, **srandom**, **initstate**, or **setstate** (“random, srandom, initstate, or setstate Subroutine” on page 7) subroutine.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## rand\_r Subroutine

### Purpose

Generates pseudo-random numbers.

### Libraries

Thread-Safe C Library (**libc\_r.a**)

Berkeley Compatibility Library (**libbsd.a**)

### Syntax

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

```
int rand_r (Seed)
unsigned int * Seed;
```

## Description

The **rand\_r** subroutine generates and returns a pseudo-random number using a multiplicative congruential algorithm. The random-number generator has a period of  $2^{32}$ , and it returns successive pseudo-random numbers.

**Note:** The **rand\_r** subroutine is a simple random-number generator. Its spectral properties (the mathematical measurement of the randomness of a number sequence) are limited. See the **drand48** subroutine or the **random** (“random, srandom, initstate, or setstate Subroutine”) subroutine for more elaborate random-number generators that have greater spectral properties.

Programs using this subroutine must link to the **libpthreads.a** library.

## Parameter

*Seed* Specifies an initial seed value.

## Return Values

0 Indicates that the subroutines was successful.  
-1 Indicates that the subroutines was not successful.

## Error Codes

If the following condition occurs, the **rand\_r** subroutine sets the **errno** global variable to the corresponding value.

**EINVAL** The *Seed* parameter specifies a null value.

## File

**/usr/include/sys/types.h** Defines system macros, data types, and subroutines.

## Related Information

The **drand48** subroutine, **random** (“random, srandom, initstate, or setstate Subroutine”) subroutine.

Subroutines Overview and List of Multithread Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## random, srandom, initstate, or setstate Subroutine

### Purpose

Generates pseudo-random numbers more efficiently.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <stdlib.h>
long random ( )
void srandom (Seed)
unsigned int Seed;

char *initstate ( Seed, State, Number)
unsigned int Seed;
char *State;
size_t Number;
char *setstate (State)
const char *State;
```

## Description

**Attention:** Do not use the **random**, **srandom**, **initstate**, or **setstate** subroutine in a multithreaded environment.

The **random** subroutine uses a non-linear 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 very large, approximately  $16 * (2^{31}-1)$ . The size of the state array determines the period of the random number generator. Increasing the state array size increases the period.

With a full 256 bytes of state information, the period of the random-number generator is greater than  $2^{69}$ , which should be sufficient for most purposes.

The **random** and **srandom** subroutines have almost the same calling sequence and initialization properties as the **rand** and **srand** subroutines. The difference is that the **rand** subroutine produces a much less random sequence; in fact, the low dozen bits generated by the **rand** subroutine go through a cyclic pattern. All the bits generated by the **random** subroutine are usable. For example, `random( ) & 01` produces a random binary value.

The **srandom** subroutine, unlike the **srand** subroutine, does not return the old seed because the amount of state information used is more than a single word. The **initstate** subroutine and **setstate** subroutine handle restarting and changing random-number generators. Like the **rand** subroutine, however, the **random** subroutine by default produces a sequence of numbers that can be duplicated by calling the **srandom** subroutine with 1 as the seed.

The **initstate** subroutine allows a state array, passed in as an argument, to be initialized for future use. The size of the state array (in bytes) is used by the **initstate** subroutine, to decide how sophisticated a random-number generator it should use; the larger the state array, the more random are the numbers. Values for the amount of state information are 8, 32, 64, 128, and 256 bytes. For amounts greater than or equal to 8 bytes, or less than 32 bytes, the **random** subroutine uses a simple linear congruential random number generator, while other amounts are rounded down to the nearest known value. The *Seed* parameter specifies a starting point for the random-number sequence and provides for restarting at the same point. The **initstate** subroutine returns a pointer to the previous state information array.

Once a state has been initialized, the **setstate** subroutine allows rapid switching between states. The array defined by *State* parameter is used for further random-number generation until the **initstate** subroutine is called or the **setstate** subroutine is called again. The **setstate** subroutine returns a pointer to the previous state array.

After initialization, a state array can be restarted at a different point in one of two ways:

- The **initstate** subroutine can be used, with the desired seed, state array, and size of the array.

- The **setstate** subroutine, with the desired state, can be used, followed by the **srandom** subroutine with the desired seed. The advantage of using both of these subroutines is that the size of the state array does not have to be saved once it is initialized.

## Parameters

<i>Seed</i>	Specifies an initial seed value.
<i>State</i>	Points to the array of state information.
<i>Number</i>	Specifies the size of the state information array.

## Error Codes

If the **initstate** subroutine is called with less than 8 bytes of state information, or if the **setstate** subroutine detects that the state information has been damaged, error messages are sent to standard error.

## Related Information

The **drand48**, **erand48**, **jrand48**, **lcong48**, **lrand48**, **mrnd48**, **nrnd48**, **seed48**, or **srand48** subroutine, **rand** or **srand** (“rand or srand Subroutine” on page 5) subroutine.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

## ra\_attachrset Subroutine

### Purpose

Attaches a work component to a resource set.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
int ra_attachrset (rstype, rsid, rset, flags)
rstype_t rstype;
rsid_t rsid;
rsethandle_t rset;
unsigned int flags;
```

### Description

The **ra\_attachrset** subroutine attaches a work component specified by the *rstype* and *rsid* parameters to a resource set specified by the *rset* parameter.

The work component is an existing process identified by the process ID. A process ID value of **RS\_MYSELF** indicates the attachment applies to the current process.

The following conditions must be met to successfully attach a process to a resource set:

- The resource set must contain processors that are available in the system.
- The calling process must either have root authority or have **CAP\_NUMA\_ATTACH** capability.
- The calling process must either have root authority or the same effective userid as the target process.
- The target process must not contain any threads that have bindprocessor bindings to a processor.
- The resource set must be contained in (be a subset of ) the target process’ partition resource set.

If any of these conditions are not met, the attachment will fail.

Once a process is attached to a resource set, the threads in the process will only run on processors contained in the resource set.

Dynamic Processor Deallocation and DLPAR may invalidate the processor attachment that is being specified. A program must become DLPAR Aware to resolve this problem.

## Parameters

<i>rstype</i>	Specifies the type of work component to be attached to the resource set specified by the <i>rset</i> parameter. The <i>rstype</i> parameter must be the following value, defined in <b>rset.h</b> : <ul style="list-style-type: none"><li>• R_PROCESS: existing process</li></ul>
<i>rsid</i>	Identifies the work component to be attached to the resource set specified by the <i>rset</i> parameter. The <i>rsid</i> parameter must be the following: <ul style="list-style-type: none"><li>• Process ID (for <i>rstype</i> of R_PROCESS): set the <i>rsid_t at_pid</i> field to the desired process' process ID.</li></ul>
<i>rset</i>	Specifies which work component (specified by the <i>rstype</i> and <i>rsid</i> parameters) to attach to the resource set.
<i>flags</i>	Reserved for future use. Specify as 0.

## Return Values

If successful, a value of 0 is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **ra\_attachrset** subroutine is unsuccessful if one or more of the following are true:

<b>EINVAL</b>	One of the following is true: <ul style="list-style-type: none"><li>• The <i>flags</i> parameter contains an invalid value.</li><li>• The <i>rstype</i> parameter contains an invalid type qualifier.</li></ul>
<b>ENODEV</b>	The resource set specified by the <i>rset</i> parameter does not contain any available processors.
<b>ESRCH</b>	The process identified by the <i>rstype</i> and <i>rsid</i> parameters does not exist.
<b>EPERM</b>	One of the following is true: <ul style="list-style-type: none"><li>• The resource set specified by the <i>rset</i> parameter is not included in the partition resource set of the process identified by the <i>rstype</i> and <i>rsid</i> parameters.</li><li>• The calling process has neither root authority nor CAP_NUMA_ATTACH attachment privilege.</li><li>• The calling process has neither root authority nor the same effective user ID as the process identified by the <i>rstype</i> and <i>rsid</i> parameters.</li><li>• The process identified by the <i>rstype</i> and <i>rsid</i> parameters has one or more threads with a bindprocessor processor binding.</li></ul>

## Related Information

“ra\_fork Subroutine” on page 14, “ra\_exec Subroutine” on page 12, “ra\_getrset Subroutine” on page 15, and “ra\_detachrset Subroutine” on page 11.

The Dynamic Logical Partitioning article in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

The **dr\_reconfig** system call in *AIX 5L Version 5.2 Technical Reference: Kernel and Subsystems Volume 1*.

---

## ra\_detachrset Subroutine

### Purpose

Detaches a work component from a resource set.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
int ra_detachrset (rstype, rsid, flags)
rstype_t rstype;
rsid_t rsid;
unsigned int flags;
```

### Description

The **ra\_detachrset** subroutine detaches a work component specified by *rstype* and *rsid* from a resource set.

The work component is an existing process identified by the process ID. A process ID value of **RS\_MYSELF** indicates the detach command applies to the current process.

The following conditions must be met to detach a process from a resource set:

- The calling process must either have root authority or have **CAP\_NUMA\_ATTACH** capability.
- The calling process must either have root authority or the same effective userid as the target process.

If these conditions are not met, the operation will fail.

Once a process is detached from a resource set, the threads in the process can run on all available processors contained in the process' partition resource set.

### Parameters

<i>rstype</i>	Specifies the type of work component to be detached from to the resource set specified by <i>rset</i> . This parameter must be the following value, defined in <b>rset.h</b> : <ul style="list-style-type: none"><li>• <b>R_PROCESS</b>: existing process</li></ul>
<i>rsid</i>	Identifies the work component to be attached to the resource set specified by <i>rset</i> . This parameter must be the following: <ul style="list-style-type: none"><li>• Process ID (for <i>rstype</i> of <b>R_PROCESS</b>): set the <i>rsid_t at_pid</i> field to the desired process' process ID.</li></ul>
<i>flags</i>	Reserved for future use. Specify as 0.

### Return Values

If successful, a value of 0 is returned. If unsuccessful, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

### Error Codes

The **ra\_detachrset** subroutine is unsuccessful if one or more of the following are true:

<b>EINVAL</b>	One of the following is true: <ul style="list-style-type: none"><li>• The <i>flags</i> parameter contains an invalid value.</li><li>• The <i>rstype</i> parameter contains an invalid type qualifier.</li></ul>
---------------	---

<b>ESRCH</b>	The process identified by the <i>rstype</i> and <i>rsid</i> parameters does not exist.
<b>EPERM</b>	One of the following is true: <ul style="list-style-type: none"> <li>• The calling process has neither root authority nor CAP_NUMA_ATTACH attachment privilege.</li> <li>• The calling process has neither root authority nor the same effective user ID as the process identified by the <i>rstype</i> and <i>rsid</i> parameters.</li> </ul>

## Related Information

“*ra\_fork* Subroutine” on page 14, “*ra\_exec* Subroutine”, “*ra\_getrset* Subroutine” on page 15, and “*ra\_attachrset* Subroutine” on page 9.

---

## *ra\_exec* Subroutine

### Purpose

Executes a file and attaches it to a given resource.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
int ra_execl(rstype, rsid, flags, path, argument0 [,argument1,...], 0)
rstype_t rstype;
rsid_t rsid;
unsigned int flags;
const char * path, argument0, argument1,...;

int ra_execle(rstype, rsid, flags, path, argument0[,argument1,...], 0, envptr)
rstype_t rstype;
rsid_t rsid;
unsigned int flags;
const char * path, argument0, argument1,...;
char * const envptr[];

int ra_execlp(rstype, rsid, flags, File, argument0[,argument1,...], 0)
rstype_t rstype;
rsid_t rsid;
unsigned int flags;
const char * File, argument0, argument1,...;

int ra_execv (rstype, rsid, flags, path, argumentv)
rstype_t rstype;
rsid_t rsid;
unsigned int flags;
const char * path;
char * const argumentv[];

int ra_execve (rstype, rsid, flags, path, argumentv, envptr)
rstype_t rstype;
rsid_t rsid;
unsigned int flags;
const char * path;
char * const argumentv[], envptr[];

int ra_execvp (rstype, rsid, flags, File, argumentv)
rstype_t rstype;
rsid_t rsid;
unsigned int flags;
const char * File;
char * const argumentv[];
```

```
int ra_exec(rstype, rsid, flags, path, argumentv, envptr)
rstype_t rstype;
rsid_t rsid;
unsigned int flags;
char * path, argumentv, envptr[];
```

## Description

The **ra\_exec** subroutine in all its forms, executes a new program in the calling process, and attaches the process to the resource specified by the *rstype* and *rsid* parameters.

The following conditions must be met to successfully attach a process to a resource set:

- The resource set must contain processors that are available in the system.
- The process must either have root authority or have CAP\_NUMA\_ATTACH capability.
- The calling thread must not have a bindprocessor binding to a processor.
- The resource set must be contained in (be a subset of ) the process' partition resource set.

**Note:** When the **exec** subroutine is used, the new process image inherits its process' resource set attachments.

Dynamic Processor Deallocation and DLPAR may invalidate the processor attachment that is being specified. A program must become DLPAR Aware to resolve this problem.

## Parameters

The **ra\_exec** subroutine has the same parameters as the **exec** subroutine, with the addition of the following new parameters:

<i>rstype</i>	Specifies the type of resource the new process image will be attached to. This parameter must be the following, defined in <b>rset.h</b> : <ul style="list-style-type: none"> <li>• R_RSET: resource set</li> </ul>
<i>rsid</i>	Identifies the resource the new process image will be attached to. This parameter must be a resource set handle. <ul style="list-style-type: none"> <li>• Process ID (for <i>rstype</i> of R_PROCESS): set the <i>rsid_t at_pid</i> field to the desired process' process ID.</li> </ul>
<i>flags</i>	Reserved for future use. Specify as 0.

## Return Values

The **ra\_exec** subroutine's return values are the same as the **exec** subroutine's return values.

## Error Codes

The **ra\_exec** subroutine's error codes are the same as the **exec** subroutine's error codes, with the addition of the following error codes:

<b>EINVAL</b>	One of the following is true: <ul style="list-style-type: none"> <li>• The <i>rstype</i> parameter contains an invalid type identifier.</li> <li>• The <i>flags</i> parameter contains an invalid flags value.</li> </ul>
<b>ENODEV</b>	The specified resource set does not contain any available processors.
<b>EFAULT</b>	Invalid address.
<b>EPERM</b>	One of the following is true: <ul style="list-style-type: none"> <li>• The calling process has neither root authority nor CAP_NUMA_ATTACH attachment privilege.</li> <li>• The calling process contains one or more threads with a bindprocessor processor binding.</li> <li>• The specified resource set is not included in the calling process' partition resource set.</li> </ul>

## Related Information

The “`ra_fork` Subroutine”, “`ra_attachrset` Subroutine” on page 9, “`ra_detachrset` Subroutine” on page 11, and “`ra_getrset` Subroutine” on page 15.

The Dynamic Logical Partitioning article in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

The `dr_reconfig` system call in *AIX 5L Version 5.2 Technical Reference: Kernel and Subsystems Volume 1*.

The `exec: execl, execl, execlp, execv, execve, execvp, or exect` Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

---

## ra\_fork Subroutine

### Purpose

Creates and attaches a new process to a given resource.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
pid_t ra_fork(rstype, rsid, flags)
rstype_t rstype;
rsid_t rsid;
unsigned int flags;
```

### Description

The `ra_fork` subroutine creates a new process, and attaches the new process to the resource set specified by `rstype` and `rsid`.

The following conditions must be met to successfully attach a process to a resource set:

- The resource set must contain processors that are available in the system.
- The process must either have root authority or have `CAP_NUMA_ATTACH` capability.
- The calling thread must not have a `bindprocessor` binding to a processor.
- The resource set must be contained in (be a subset of ) the process' partition resource set.

**Note:** When the `fork` subroutine is used, the child process inherits its parent's resource set attachments.

Dynamic Processor Deallocation and DLPAR may invalidate the processor attachment that is being specified. A program must become DLPAR Aware to resolve this problem.

### Parameters

<i>rstype</i>	Specifies the type of resource the new process will be attached to. This parameter must be the following value, defined in <b>rset.h</b> . <ul style="list-style-type: none"><li>• <code>R_RSET</code>: resource set.</li></ul>
<i>rsid</i>	Identifies the resource the new process will be attached to. This parameter must be a resource set handle. <ul style="list-style-type: none"><li>• Process ID (for <i>rstype</i> of <code>R_PROCESS</code>): set the <i>rsid_t at_pid</i> field to the desired process' process ID.</li></ul>
<i>flags</i>	Reserved for future use. Specify as 0.

## Return Values

The `ra_fork` subroutine's return values are the same as the `fork` subroutine's return values.

## Error Codes

The `ra_fork` subroutine's error codes are the same as the `fork` subroutine's error codes with the addition of the following:

<b>EINVAL</b>	One of the following is true: <ul style="list-style-type: none"><li>• The <code>rstype</code> parameter contains an invalid type identifier.</li><li>• The <code>flags</code> parameter contains an invalid flags value.</li></ul>
<b>ENODEV</b>	The specified resource set does not contain any available processors.
<b>EFAULT</b>	Invalid address.
<b>EPERM</b>	One of the following is true: <ul style="list-style-type: none"><li>• The calling process has neither root authority nor <code>CAP_NUMA_ATTACH</code> attachment privilege.</li><li>• The calling process contains one or more threads with a bindprocessor processor binding.</li><li>• The specified resource set is not included in the calling process' partition resource set.</li></ul>

## Related Information

“`ra_attachrset` Subroutine” on page 9, “`ra_detachrset` Subroutine” on page 11, and “`ra_getrset` Subroutine”.

The Dynamic Logical Partitioning article in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

The `dr_reconfig` system call in *AIX 5L Version 5.2 Technical Reference: Kernel and Subsystems Volume 1*.

The `fork`, `f_fork`, or `vfork` Subroutine, and `exec`: `execl`, `execle`, `execlp`, `execv`, `execve`, `execvp`, or `exec` Subroutine articles in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 2*.

---

## ra\_getrset Subroutine

### Purpose

Gets the resource set to which a work component is attached.

### Library

Standard C library (`libc.a`)

### Syntax

```
# include <sys/rset.h>
int ra_getrset (rstype, rsid, flags, rset)
rstype_t rstype;
rsid_t rsid;
unsigned int flags;
rsethandle_t rset;
```

### Description

The `ra_getrset` subroutine returns the resource set to which a specified work component is attached.

The work component is an existing process identified by the process ID. A process ID value of `RS_MYSELF` indicates the resource set attached to the current process is requested.

The following return values from the **ra\_getrset** subroutine indicate the type of resource set returned:

- A value of RS\_EFFECTIVE\_RSET indicates the process was explicitly attached to the resource set. This may have been done with the **ra\_attachrset** subroutine.
- A value of RS\_PARTITION\_RSET indicates the process was not explicitly attached to a resource set. However, the process had an explicitly set partition resource set. This may be set with the **rs\_setpartition** subroutine or through the use of WLM work classes with resource sets.
- A value of RS\_DEFAULT\_RSET indicates the process was not explicitly attached to a resource set nor did it have an explicitly set partition resource set. The system default resource set is returned.

## Parameters

<i>rstype</i>	Specifies the type of the work component whose resource set attachment is requested. This parameter must be the following value, defined in <b>rset.h</b> : <ul style="list-style-type: none"><li>• R_PROCESS: existing process</li></ul>
<i>rsid</i>	Identifies the work component whose resource set attachment is requested. This parameter must be the following: <ul style="list-style-type: none"><li>• Process ID (for <i>rstype</i> of R_PROCESS): set the <i>rsid_t at_pid</i> field to the desired process' process ID.</li></ul>
<i>flags</i>	Reserved for future use. Specify as 0.
<i>rset</i>	Specifies the resource set to receive the work component's resource set.

## Return Values

If successful, a value of RS\_EFFECTIVE\_RSET, RS\_PARTITION\_RSET, or RS\_DEFAULT\_RSET is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **ra\_getrset** subroutine is unsuccessful if one or more of the following are true:

<b>EINVAL</b>	One of the following is true: <ul style="list-style-type: none"><li>• The <i>flags</i> parameter contains an invalid value.</li><li>• The <i>rstype</i> parameter contains an invalid type qualifier.</li></ul>
<b>EFAULT</b>	Invalid address.
<b>ESRCH</b>	The process identified by the <i>rstype</i> and <i>rsid</i> parameters does not exist.

## Related Information

The “rs\_getpartition Subroutine” on page 93.

---

## read, readx, readv, readvx, or pread Subroutine

### Purpose

Reads from a file.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
ssize_t read (FileDescriptor,  
             Buffer, NBytes)
```

```

int FileDescriptor;
void * Buffer;
size_t NBytes;

int readx (FileDescriptor, ;Buffer, NBytes, Extension)
int FileDescriptor;
char * Buffer;
unsigned int NBytes;
int Extension;
#include <sys/uio.h>

ssize_t readv (FileDescriptor, iov, iovCount)
int FileDescriptor;
const struct iovec * iov;
int iovCount;

ssize_t readvx (FileDescriptor, iov, iovCount, Extension)
int FileDescriptor;
struct iovec *iov;
int iovCount;
int Extension;
#include <unistd.h>
ssize_t pread (int fildes, void *buf, size_t nbyte, off_t offset);

```

## Description

The **read** subroutine attempts to read *NBytes* of data from the file associated with the *FileDescriptor* parameter into the buffer pointed to by the *Buffer* parameter.

The **readv** subroutine performs the same action but scatters the input data into the *iovCount* buffers specified by the array of **iovec** structures pointed to by the *iov* parameter. Each **iovec** entry specifies the base address and length of an area in memory where data should be placed. The **readv** subroutine always fills an area completely before proceeding to the next.

The **readx** and **readvx** subroutines are the same as the **read** and **readv** subroutines, respectively, with the addition of an *Extension* parameter, which is needed when reading from some device drivers and when reading directories. While directories can be read directly, it is recommended that the **opendir** and **readdir** calls be used instead, as this is a more portable interface.

On regular files and devices capable of seeking, the **read** starts at a position in the file given by the file pointer associated with the *FileDescriptor* parameter. Upon return from the **read** subroutine, the file pointer is incremented by the number of bytes actually read.

Devices that are incapable of seeking always read from the current position. The value of a file pointer associated with such a file is undefined.

On directories, the **readvx** subroutine starts at the position specified by the file pointer associated with the *FileDescriptor* parameter. The value of this file pointer must be either 0 or a value which the file pointer had immediately after a previous call to the **readvx** subroutine on this directory. Upon return from the **readvx** subroutine, the file pointer increments by a number that may not correspond to the number of bytes copied into the buffers.

When attempting to read from an empty pipe (first-in-first-out (FIFO)):

- If no process has the pipe open for writing, the **read** returns 0 to indicate end-of-file.
- If some process has the pipe open for writing:
  - If **O\_NDELAY** and **O\_NONBLOCK** are clear (the default), the **read** blocks until some data is written or the pipe is closed by all processes that had opened the pipe for writing.

- If **O\_NDELAY** is set, the **read** subroutine returns a value of 0.
- If **O\_NONBLOCK** is set, the **read** subroutine returns a value of -1 and sets the global variable **errno** to **EAGAIN**.

When attempting to read from a character special file that supports nonblocking reads, such as a terminal, and no data is currently available:

- If **O\_NDELAY** and **O\_NONBLOCK** are clear (the default), the **read** subroutine blocks until data becomes available.
- If **O\_NDELAY** is set, the **read** subroutine returns 0.
- If **O\_NONBLOCK** is set, the **read** subroutine returns -1 and sets the **errno** global variable to **EAGAIN** if no data is available.

When attempting to read a regular file that supports enforcement mode record locks, and all or part of the region to be read is currently locked by another process:

- If **O\_NDELAY** and **O\_NONBLOCK** are clear, the **read** blocks the calling process until the lock is released.
- If **O\_NDELAY** or **O\_NONBLOCK** is set, the **read** returns -1 and sets the global variable **errno** to **EAGAIN**.

The behavior of an interrupted **read** subroutine depends on how the handler for the arriving signal was installed.

**Note:** A read from a regular file is not interruptible. Only reads from objects that may block indefinitely, such as FIFOs, sockets, and some devices, are generally interruptible.

If the handler was installed with an indication that subroutines should not be restarted, the **read** subroutine returns a value of -1 and the global variable **errno** is set to **EINTR** (even if some data was already consumed).

If the handler was installed with an indication that subroutines should be restarted:

- If no data had been read when the interrupt was handled, this **read** will not return a value (it is restarted).
- If data had been read when the interrupt was handled, this **read** subroutine returns the amount of data consumed.

The **pread** function performs the same action as **read**, except that it reads from a given position in the file without changing the file pointer. The first three arguments to **pread** are the same as **read** with the addition of a fourth argument offset for the desired position inside the file. An attempt to perform a **pread** on a file that is incapable of seeking results in an error.

**Note:** The **pread64** subroutine applies to AIX 4.3 and later.

```
ssize_t pread64(int filedes , void *buf , size_t nbytes , off64_t offset)
```

The **pread64** subroutines performs the same action as **pread** but the limit of offset to the maximum file size for the file associated with the file Descriptor and **DEV\_OFF\_MAX** if the file associated with fileDescriptor is a block special or character special file.

## Parameters

*FileDescriptor*                      A file descriptor identifying the object to be read.

## Extension

Provides communication with character device drivers that require additional information or return additional status. Each driver interprets the *Extension* parameter in a device-dependent way, either as a value or as a pointer to a communication area. Drivers must apply reasonable defaults when the value of the *Extension* parameter is 0.

For directories, the *Extension* parameter determines the format in which directory entries should be returned:

- If the value of the *Extension* parameter is 0, the format in which directory entries are returned depends on the value of the **real directory read** flag (described in the **ulimit** (“ulimit Subroutine” on page 415) subroutine).
- If the calling process does not have the **real directory read** flag set, the buffers are filled with an array of directory entries truncated to fit the format of the System V directory structure. This provides compatibility with programs written for UNIX System V.
- If the calling process has the **real directory read** flag set (see the **ulimit** subroutine), the buffers are filled with an image of the underlying implementation of the directory.
- If the value of the *Extension* parameter is 1, the buffers are filled with consecutive directory entries in the format of **adirent** structure. This is logically equivalent to the **readdir** subroutine.
- Other values of the *Extension* parameter are reserved.

For tape devices, the *Extension* parameter determines the response of the **readx** subroutine when the tape drive is in variable block mode and the read request is for less than the tape’s block size.

- If the value of the *Extension* parameter is **TAPE\_SHORT\_READ**, the **readx** subroutine returns the number of bytes requested and sets the **errno** global variable to a value of 0.
- If the value of the *Extension* parameter is 0, the **readx** subroutine returns a value of 0 and sets the **errno** global variable to **ENOMEM**.

## iov

Points to an array of **iovec** structures that identifies the buffers into which the data is to be placed. The **iovec** structure is defined in the **sys/uio.h** file and contains the following members:

```
caddr_t iov_base;  
size_t  iov_len;
```

*iovCount* Specifies the number of **iovec** structures pointed to by the *iov* parameter.

*Buffer* Points to the buffer.

*NBytes* Specifies the number of bytes read from the file associated with the *FileDescriptor* parameter.

**Note:** When reading tapes, the **read** subroutines consume a physical tape block on each call to the subroutine. If the physical data block size is larger than specified by the *Nbytes* parameter, an error will be returned, since all of the data from the read will not fit into the buffer specified by the read.

To avoid read errors due to unknown blocking sizes on tapes, set the *NBytes* parameter to a very large value (such as 32K bytes).

## Return Values

Upon successful completion, the **read**, **readx**, **readv**, **readvx**, and **pread** subroutines return the number of bytes actually read and placed into buffers. The system guarantees to read the number of bytes requested if the descriptor references a normal file that has the same number of bytes left before the end of the file is reached, but in no other case.

A value of 0 is returned when the end of the file has been reached. (For information about communication files, see the **ioctl** and **termio** files.)

Otherwise, a value of -1 is returned, the global variable **errno** is set to identify the error, and the content of the buffer pointed to by the *Buffer* or *iov* parameter is indeterminate.

## Error Codes

The **read**, **readx**, **readv**, **readvx**, and **pread** subroutines are unsuccessful if one or more of the following are true:

<b>EBADMSG</b>	The file is a STREAM file that is set to control-normal mode and the message waiting to be read includes a control part.
<b>EBADF</b>	The <i>FileDescriptor</i> parameter is not a valid file descriptor open for reading.
<b>EINVAL</b>	The file position pointer associated with the <i>FileDescriptor</i> parameter was negative.
<b>EINVAL</b>	The sum of the <b>iov_len</b> values in the <i>iov</i> array was negative or overflowed a 32-bit integer.
<b>EINVAL</b>	The value of the <i>iovCount</i> parameter was not between 1 and 16, inclusive.
<b>EINVAL</b>	The value of the <i>Nbytes</i> parameter that is larger than <b>OFF_MAX</b> , was requested on the 32-bit kernel. This is a case where the system call is requested from a 64-bit application that is running on a 32-bit kernel.
<b>EINVAL</b>	The STREAM or multiplexer referenced by <i>FileDescriptor</i> is linked (directly or indirectly) downstream from a multiplexer.
<b>EAGAIN</b>	The file was marked for non-blocking I/O, and no data was ready to be read.
<b>EFAULT</b>	The <i>Buffer</i> or part of the <i>iov</i> points to a location outside of the allocated address space of the process.
<b>EDEADLK</b>	A deadlock would occur if the calling process were to sleep until the region to be read was unlocked.
<b>EINTR</b>	A <b>read</b> was interrupted by a signal before any data arrived, and the signal handler was installed with an indication that subroutines are not to be restarted.
<b>EIO</b>	An I/O error occurred while reading from the file system.
<b>EIO</b>	The process is a member of a background process attempting to read from its controlling terminal, and either the process is ignoring or blocking the <b>SIGTTIN</b> signal or the process group has no parent process.
<b>EFBIG</b>	An offset greater than <b>MAX_FILESIZE</b> was requested on the 32-bit kernel.
<b>EOVERFLOW</b>	An attempt was made to read from a regular file where <i>NBytes</i> was greater than zero and the starting offset was before the end-of-file and was greater than or equal to the offset maximum established in the open file description associated with <i>FileDescriptor</i> .

The **read**, **readx**, **readv**, **readvx** and **pread** subroutines may be unsuccessful if the following is true:

<b>ENXIO</b>	A request was made of a nonexistent device, or the request was outside the capabilities of the device.
<b>ESPIPE</b>	<i>fildev</i> is associated with a pipe or FIFO.

If Network File System (NFS) is installed on the system, the **read** system call can also fail if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Related Information

The **fcntl**, **dup**, or **dup2** subroutine, **ioctl** subroutine, **lockfx** subroutine, **lseek** subroutine, **open**, **openx**, or **creat** subroutine, **opendir**, **readdir**, or **seekdir** subroutine, **pipe** subroutine, **poll** subroutine, **socket** subroutine, **socketpair** subroutine.

The Input and Output Handling in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## readdir\_r Subroutine

### Purpose

Reads a directory.

### Library

Thread-Safe C Library (**libc\_r.a**)

### Syntax

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

```
int readdir_r (DirectoryPointer, Entry, Result)
DIR * DirectoryPointer;
struct dirent * Entry;
struct dirent ** Result;
```

### Description

The **readdir\_r** subroutine returns the directory entry in the structure pointed to by the *Result* parameter. The **readdir\_r** subroutine returns entries for the . (dot) and .. (dot-dot) directories, if present, but never returns an invalid entry (with *d\_ino* set to 0). When it reaches the end of the directory, the **readdir\_r** subroutine returns a 0. When it detects an invalid **seekdir** operation, the **readdir\_r** subroutine returns a 9.

**Note:** The **readdir** subroutine is reentrant when an application program uses different *DirectoryPointer* parameter values (returned from the **opendir** subroutine). Use the **readdir\_r** subroutine when multiple threads use the same directory pointer.

Using the **readdir\_r** subroutine after the **closedir** subroutine, for the structure pointed to by the *DirectoryPointer* parameter, has an undefined result. The structure pointed to by the *DirectoryPointer* parameter becomes invalid for all threads, including the caller.

Programs using this subroutine must link to the **libpthread.a** library.

### Parameters

<i>DirectoryPointer</i>	Points to the <b>DIR</b> structure of an open directory.
<i>Entry</i>	Points to a structure that contains the next directory entry.
<i>Result</i>	Points to the directory entry specified by the <i>Entry</i> parameter.

### Return Values

**0** Indicates that the subroutines was successful.  
**9** Indicates that the subroutines was not successful.

### Error Codes

If the **readdir\_r** subroutine is unsuccessful, the **errno** global variable is set to one of the following values:

<b>EACCES</b>	Search permission is denied for any component of the structure pointed to by the <i>DirectoryPointer</i> parameter, or read permission is denied for the structure pointed to by the <i>DirectoryPointer</i> parameter.
---------------	---

<b>ENAMETOOLONG</b>	The length of the <i>DirectoryPointer</i> parameter exceeds the value of the <b>PATH_MAX</b> variable, or a path-name component is longer than the value of <b>NAME_MAX</b> variable while the <b>_POSIX_NO_TRUNC</b> variable is in effect.
<b>ENOENT</b>	The named directory does not exist.
<b>ENOTDIR</b>	A component of the structure pointed to by the <i>DirectoryPointer</i> parameter is not a directory.
<b>EMFILE</b>	Too many file descriptors are currently open for the process.
<b>ENFILE</b>	Too many file descriptors are currently open in the system.
<b>EBADF</b>	The structure pointed to by the <i>DirectoryPointer</i> parameter does not refer to an open directory stream.

## Examples

To search a directory for the entry name, enter:

```
len = strlen(name);
DirectoryPointer = opendir(".");
for (readdir_r(DirectoryPointer, &Entry, &Result); Result != NULL;
     readdir_r(DirectoryPointer, &Entry, &Result))
    if (dp->d_namlen == len && !strcmp(dp->d_name, name)) {
        closedir(DirectoryPointer);
        return FOUND;
    }
closedir(DirectoryPointer);
return NOT_FOUND;
```

## Related Information

The **close** subroutine, **exec** subroutines, **fork** subroutine, **lseek** subroutine, **openx**, **open**, or **creat** subroutine, **read**, **readv**, **readx**, or **readvx** (“read, readx, readv, readvx, or pread Subroutine” on page 16) subroutine, **scandir** or **alphasort** (“scandir or alphasort Subroutine” on page 107) subroutine.

The **opendir**, **readdir**, **telldir**, **seekdir**, **rewinddir**, or **closedir** subroutine.

Subroutines Overview, List of File and Directory Manipulation Services, and List of Multithread Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## readlink Subroutine

### Purpose

Reads the contents of a symbolic link.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
int readlink ( Path, Buffer, BufferSize)
const char *Path;
char *Buffer;
size_t BufferSize;
```

### Description

The **readlink** subroutine copies the contents of the symbolic link named by the *Path* parameter in the buffer specified in the *Buffer* parameter. The *BufferSize* parameter indicates the size of the buffer in bytes.

If the actual length of the symbolic link is less than the number of bytes specified in the *BufferSize* parameter, the string copied into the buffer will be null-terminated. If the actual length of the symbolic link is greater than the number of bytes specified in the *BufferSize* parameter, an error is returned. The length of a symbolic link cannot exceed 1023 characters or the value of the **PATH\_MAX** constant. **PATH\_MAX** is defined in the **limits.h** file.

## Parameters

<i>Path</i>	Specifies the path name of the destination file or directory.
<i>Buffer</i>	Points to the user's buffer. The buffer should be at least as large as the <i>BufferSize</i> parameter.
<i>BufferSize</i>	Indicates the size of the buffer. The contents of the link are null-terminated, provided there is room in the buffer.

## Return Values

Upon successful completion, the **readlink** subroutine returns a count of the number of characters placed in the buffer (not including any terminating null character). If the **readlink** subroutine is unsuccessful, the buffer is not modified, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **readlink** subroutine fails if one or both of the following are true:

<b>ENOENT</b>	The file named by the <i>Path</i> parameter does not exist, or the path points to an empty string.
<b>EINVAL</b>	The file named by the <i>Path</i> parameter is not a symbolic link.
<b>ERANGE</b>	The path name in the symbolic link is longer than the <i>BufferSize</i> value.

The **readlink** subroutine can also fail due to additional errors. See "Base Operating System Error Codes for Services that Require Path-Name Resolution" for a list of additional error codes.

The **readlink** subroutine can also fail due to additional errors. See Appendix A, "Base Operating System Error Codes for Services That Require Path-Name Resolution" on page A-1 for a list of additional error codes.

If Network File System (NFS) is installed on the system, the **readlink** subroutine can also fail if the following is true:

<b>ETIMEDOUT</b>	The connection timed out.
------------------	---------------------------

## Related Information

The **In** command.

The **link** subroutine, **statx**, **stat**, **fstatx**, **fstat**, **fullstat**, or **ffullstat** ("statx, stat, lstat, fstatx, fstat, fullstat, ffullstat, stat64, lstat64, or fstat64 Subroutine" on page 277) subroutine, **symlink** ("symlink Subroutine" on page 308) subroutine, **unlink** ("unlink Subroutine" on page 423) subroutine.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## read\_real\_time or time\_base\_to\_time Subroutine

### Purpose

Read the processor real time clock or time base registers to obtain high-resolution elapsed time.

## Library

Standard C Library (**libc.a**)

## Syntax

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

int read_real_time(timebasestruct_t *t,
                  size_t size_of_timebasestruct_t);

int time_base_to_time(timebasestruct_t *t,
                     size_t size_of_timebasestruct_t);
```

## Description

These subroutines are designed to be used for making high-resolution measurement of elapsed time, using the processor real time clock or time base registers. The **read\_real\_time** subroutine reads the value of the appropriate registers and stores them in a structure. The **time\_base\_to\_time** subroutine converts time base data to real time, if necessary. This process is divided into two steps because the process of reading the time is usually part of the timed code, and so the conversion from time base to real time can be moved out of the timed code.

The **read\_real\_time** subroutine reads either the processor real time clock (for the POWER family or PowerPC 601 RISC Microprocessor in AIX 5.1 and earlier) or the time base register (in the case of the POWER-based processors other than the PowerPC 601 RISC Microprocessor). The *t* argument is a pointer to a *timebasestruct\_t*, where the time values are recorded.

After calling **read\_real\_time**, if running on a processor with a real time clock, *t->tb\_high* and *t->tb\_low* contain the current clock values (seconds and nanoseconds), and *t->flag* contains the **RTC\_POWER**.

If running on a processor with a time base register, *t->tb\_high* and *t->tb\_low* contain the current values of the time base register, and *t->flag* contains **RTC\_POWER\_PC**.

The **time\_base\_to\_time** subroutine converts time base information to real time, if necessary. It is recommended that applications unconditionally call the **time\_base\_to\_time** subroutine rather than performing a check to see if it is necessary.

If *t->flag* is **RTC\_POWER**, the subroutine simply returns (the data is already in real time format).

If *t->flag* is **RTC\_POWER\_PC**, the time base information in *t->tb\_high* and *t->tb\_low* is converted to seconds and nanoseconds; *t->tb\_high* is replaced by the seconds; *t->tb\_low* is replaced by the nanoseconds; and *t->flag* is changed to **RTC\_POWER**.

## Parameters

*t* Points to a *timebasestruct\_t*.

## Return Values

The **read\_real\_time** subroutine returns **RTC\_POWER** if the contents of the real time clock has been recorded in the *timebasestruct*, or returns **RTC\_POWER\_PC** if the content of the time base registers has been recorded in the *timebasestruct*.

The **time\_base\_to\_time** subroutine returns **0** if the conversion to real time is successful (or not necessary), otherwise **-1** is returned.

## Examples

This example shows the time it takes for **printf** to print the comment between the begin and end time codes:

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

int
main(void)
{
    timebasestruct_t start, finish;
    int val = 3;
    int secs, n_secs;

    /* get the time before the operation begins */
    read_real_time(&start, TIMEBASE_SZ);

    /* begin code to be timed */
    (void) printf("This is a sample line %d \n", val);
    /* end code to be timed */

    /* get the time after the operation is complete */
    read_real_time(&finish, TIMEBASE_SZ);

    /*
     * Call the conversion routines unconditionally, to ensure
     * that both values are in seconds and nanoseconds regardless
     * of the hardware platform.
     */
    time_base_to_time(&start, TIMEBASE_SZ);
    time_base_to_time(&finish, TIMEBASE_SZ);

    /* subtract the starting time from the ending time */
    secs = finish.tb_high - start.tb_high;
    n_secs = finish.tb_low - start.tb_low;

    /*
     * If there was a carry from low-order to high-order during
     * the measurement, we may have to undo it.
     */
    if (n_secs < 0) {
        secs--;
        n_secs += 1000000000;
    }

    (void) printf("Sample time was %d seconds %d nanoseconds\n",
                 secs, n_secs);

    exit(0);
}
```

## Related Information

The **gettimer**, **settimer**, **restimer**, **stime**, or **time** subroutines, **getrusage**, **times**, or **vtimes** subroutines.

High-Resolution Time Measurements Using POWER-based Time Base or POWER family Real-Time Clock in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## realpath Subroutine

### Purpose

Resolves path names.

## Library

Standard C Library (**libc.a**)

## Syntax

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

```
char *realpath (const char *file_name,  
char *resolved_name)
```

## Description

The **realpath** subroutine performs filename expansion and path name resolution in *file\_name* and stores it in *resolved\_name*.

The **realpath** subroutine can handle both relative and absolute path names. For both absolute and relative path names, the **realpath** subroutine returns the resolved absolute path name.

The character pointed to by *resolved\_name* must be big enough to contain the fully resolved path name. The value of `PATH_MAX` (defined in **limits.h** header file) may be used as an appropriate array size.

## Return Values

On successful completion, the **realpath** subroutine returns a pointer to the resolved name. Otherwise, it returns a null pointer, and sets **errno** to indicate the error. If the **realpath** subroutine encounters an error, the contents of *resolved\_name* are undefined.

## Error Codes

Under the following conditions, the **realpath** subroutine fails and sets **errno** to:

<b>EACCES</b>	Read or search permission was denied for a component of the path name.
<b>EINVAL</b>	<i>file_name</i> or <i>resolved_name</i> is a null pointer.
<b>ELOOP</b>	Too many symbolic links are encountered in translating <i>file_name</i> .
<b>ENAMETOOLONG</b>	The length of <i>file_name</i> or <i>resolved_name</i> exceeds <code>PATH_MAX</code> or a path name component is longer than <code>NAME_MAX</code> .
<b>ENOENT</b>	The <i>file_name</i> parameter does not exist or points to an empty string.
<b>ENOTDIR</b>	A component of the <i>file_name</i> prefix is not a directory.

The **realpath** subroutine may fail if:

**ENOMEM** Insufficient storage space is available.

## Related Information

The **getcwd** or **sysconf** (“**sysconf** Subroutine” on page 311) subroutine.

---

## reboot Subroutine

### Purpose

Restarts the system.

## Library

Standard C Library (**libc.a**)

## Syntax

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

```
void reboot ( HowTo, Argument )  
int HowTo;  
void *Argument;
```

## Description

The **reboot** subroutine restarts or re-initial program loads (IPL) the system. The startup is automatic and brings up **/unix** in the normal, nonmaintenance mode.

**Note:** The routine may coredump instead of returning EFAULT when an invalid pointer is passed in case of 64-bit application calling 32-bit kernel interface.

The calling process must have root user authority in order to run this subroutine successfully.

**Attention:** Users of the **reboot** subroutine are not portable. The **reboot** subroutine is intended for use only by the **halt**, **reboot**, and **shutdown** commands.

## Parameters

*HowTo* Specifies one of the following values:

**RB\_SOFTIPL**

Soft IPL.

**RB\_HALT**

Halt operator; turn the power off.

**RB\_POWIPL**

Halt operator; turn the power off. Wait a specified length of time, and then turn the power on.

*Argument* Specifies the amount of time (in seconds) to wait between turning the power off and turning the power on. This option is not supported on all models. Please consult your hardware technical reference for more details.

## Return Values

Upon successful completion, the **reboot** subroutine does not return a value. If the **reboot** subroutine fails, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **reboot** subroutine is unsuccessful if any of the following is true:

**EPERM**

The calling process does not have root user authority.

**EINVAL**

The *HowTo* value is not valid.

**EFAULT**

The *Argument* value is not a valid address.

## Related Information

The **halt** command, **reboot** command, **shutdown** command.

---

## re\_comp or re\_exec Subroutine

### Purpose

Regular expression handler.

### Library

Standard C Library (**libc.a**)

### Syntax

```
char *re_comp( String)
const char *String;
int re_exec(String)
const char *String;
```

### Description

**Attention:** Do not use the **re\_comp** or **re\_exec** subroutine in a multithreaded environment.

The **re\_comp** subroutine compiles a string into an internal form suitable for pattern matching. The **re\_exec** subroutine checks the argument string against the last string passed to the **re\_comp** subroutine.

The **re\_comp** subroutine returns 0 if the string pointed to by the *String* parameter was compiled successfully; otherwise a string containing an error message is returned. If the **re\_comp** subroutine is passed 0 or a null string, it returns without changing the currently compiled regular expression.

The **re\_exec** subroutine returns 1 if the string pointed to by the *String* parameter matches the last compiled regular expression, 0 if the string pointed to by the *String* parameter failed to match the last compiled regular expression, and -1 if the compiled regular expression was invalid (indicating an internal error).

The strings passed to both **re\_comp** and **re\_exec** subroutines may have trailing or embedded newline characters; they are terminated by nulls. The regular expressions recognized are described in the manual entry for the **ed** command, given the above difference.

### Parameters

*String*        Points to a string that is to be matched or compiled.

### Return Values

If an error occurs, the **re\_exec** subroutine returns a -1, while the **re\_comp** subroutine returns one of the following strings:

- No previous regular expression
- Regular expression too long
- unmatched \ (
- missing ]
- too many \(\) pairs
- unmatched \)

## Related Information

The **compile**, **step**, or **advance** subroutine, **regcmp** or **regex** (“regcmp or regex Subroutine”) subroutine.

The **ed** command, **sed** command, **grep** command.

List of String Manipulation Services and Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## regcmp or regex Subroutine

### Purpose

Compiles and matches regular-expression patterns.

### Libraries

Standard C Library ( **libc.a** )

Programmers Workbench Library (**libPW.a**)

### Syntax

```
#include <libgen.h>
```

```
char *regcmp ( String [, String, . . . ], (char *) 0)  
const char *String, . . . ;
```

```
const char *regex ( Pattern, Subject [, ret, . . . ])  
char *Pattern, *Subject, *ret, . . . ;  
extern char *__loc1;
```

### Description

The **regcmp** subroutine compiles a regular expression (or *Pattern*) and returns a pointer to the compiled form. The **regcmp** subroutine allows multiple *String* parameters. If more than one *String* parameter is given, then the **regcmp** subroutine treats them as if they were concatenated together. It returns a null pointer if it encounters an incorrect parameter.

You can use the **regcmp** command to compile regular expressions into your C program, frequently eliminating the need to call the **regcmp** subroutine at run time.

The **regex** subroutine compares a compiled *Pattern* to the *Subject* string. Additional parameters are used to receive values. Upon successful completion, the **regex** subroutine returns a pointer to the next unmatched character. If the **regex** subroutine fails, a null pointer is returned. A global character pointer, **\_\_loc1**, points to where the match began.

The **regcmp** and **regex** subroutines are borrowed from the **ed** command; however, the syntax and semantics have been changed slightly. You can use the following symbols with the **regcmp** and **regex** subroutines:

[ ] \* . ^

These symbols have the same meaning as they do in the **ed** command.

-	The minus sign (or hyphen) within brackets used with the <b>regex</b> subroutine means "through," according to the current collating sequence. For example, [a-z] can be equivalent to [abcd . . . xyz] or [aBbCc . . . xYyZz]. You can use the - by itself if the - is the last or first character. For example, the character class expression [ ] -] matches the ] (right bracket) and - (minus) characters.
\$	Matches the end of the string. Use the \n character to match a new-line character.
+	A regular expression followed by + (plus sign) means one or more times. For example, [0-9] + is equivalent to [0-9] [0-9] *.
[ m ] [ m, ] [ m, u ]	Integer values enclosed in [ ] (braces) indicate the number of times to apply the preceding regular expression. The <i>m</i> character is the minimum number and the <i>u</i> character is the maximum number. The <i>u</i> character must be less than 256. If you specify only <i>m</i> , it indicates the exact number of times to apply the regular expression. [m,] is equivalent to [m,u.] and matches <i>m</i> or more occurrences of the expression. The + (plus sign) and * (asterisk) operations are equivalent to [1,] and [0,], respectively.
( . . . )\$n	This stores the value matched by the enclosed regular expression in the ( <i>n</i> +1)th <i>ret</i> parameter. Ten enclosed regular expressions are allowed. The <b>regex</b> subroutine makes the assignments unconditionally.
( . . . )	Parentheses group subexpressions. An operator, such as *, +, or [ ] works on a single character or on a regular expression enclosed in parentheses. For example, (a*(cb+*))\$0.

All of the preceding defined symbols are special. You must precede them with a \ (backslash) if you want to match the special symbol itself. For example, \\$ matches a dollar sign.

**Note:** The **regcmp** subroutine uses the **malloc** subroutine to make the space for the vector. Always free the vectors that are not required. If you do not free the unneeded vectors, you can run out of memory if the **regcmp** subroutine is called repeatedly. Use the following as a replacement for the **malloc** subroutine to reuse the same vector, thus saving time and space:

```
/* . . . Your Program . . . */
malloc(n)
{
    int n;
    static int rebuf[256] ;

    return ((n <= sizeof(rebuf)) ? rebuf : NULL);
}
```

The **regcmp** subroutine produces code values that the **regex** subroutine can interpret as the regular expression. For instance, [a-z] indicates a range expression which the **regcmp** subroutine compiles into a string containing the two end points (a and z).

The **regex** subroutine interprets the range statement according to the current collating sequence. The expression [a-z] can be equivalent either to [abcd . . . xyz] , or to [aBbCcDd . . . xXyYzZ], as long as the character *preceding* the minus sign has a lower collating value than the character *following* the minus sign.

The behavior of a range expression is dependent on the collation sequence. If you want to match a *specific* set of characters, you should list each one. For example, to select letters a, b, or c, use [abc] rather than [a-c] .

## Notes:

1. No assumptions are made at compile time about the actual characters contained in the range.
2. Do not use multibyte characters.
3. You can use the ] (right bracket) itself within a pair of brackets if it immediately follows the leading [ (left bracket) or [^ (a left bracket followed immediately by a circumflex).
4. You can also use the minus sign (or hyphen) if it is the first or last character in the expression. For example, the expression [ ] -0] matches either the right bracket ( ] ), or the characters - through 0.

## Matching a Character Class in National Language Support

A common use of the range expression is matching a character class. For example, [0-9] represents all digits, and [a-z, A-Z] represents all letters. This form may produce unexpected results when ranges are interpreted according to the current collating sequence.

Instead of the range expression shown above, use a character class expression within brackets to match characters. The system interprets this type of expression according to the current character class definition. However, you cannot use character class expressions in range expressions.

The following exemplifies the syntax of a character class expression:

```
[ :charclass:]
```

that is, a left bracket followed by a colon, followed by the name of the character class, followed by another colon and a right bracket.

National Language Support supports the following character classes:

<b>[ :upper:]</b>	ASCII uppercase letters.
<b>[ :lower:]</b>	ASCII lowercase letters.
<b>[ :alpha:]</b>	ASCII uppercase and lowercase letters.
<b>[ :digit:]</b>	ASCII digits.
<b>[ :alnum:]</b>	ASCII uppercase and lowercase letters, and digits.
<b>[ :xdigit:]</b>	ASCII hexadecimal digits.
<b>[ :punct:]</b>	ASCII punctuation character (neither a control character nor an alphanumeric character).
<b>[ :space:]</b>	ASCII space, tab, carriage return, new-line, vertical tab, or form feed character.
<b>[ :print:]</b>	ASCII printing characters.

## Parameters

<i>Subject</i>	Specifies a comparison string.
<i>String</i>	Specifies the <i>Pattern</i> to be compiled.
<i>Pattern</i>	Specifies the expression to be compared.
<i>ret</i>	Points to an address at which to store comparison data. The <b>regex</b> subroutine allows multiple <b>ret</b> <i>String</i> parameters.

## Related Information

The **ctype** subroutine, **compile**, **step**, or **advance** subroutine, **malloc**, **free**, **realloc**, **calloc**, **mallopt**, **mallinfo**, or **alloca** subroutine, **regcomp** (“regcomp Subroutine” on page 32) subroutine, **regex** (“regexec Subroutine” on page 35) subroutine.

The **ed** command, **regcmp** command.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## regcomp Subroutine

### Purpose

Compiles a specified basic or extended regular expression into an executable string.

### Library

Standard C Library (**libc. a**)

### Syntax

```
#include <regex.h>
```

```
int regcomp ( Preg, Pattern, CFlags)
const char *Preg;
const char *Pattern;
int CFlags;
```

### Description

The **regcomp** subroutine compiles the basic or extended regular expression specified by the *Pattern* parameter and places the output in the structure pointed to by the *Preg* parameter.

### Parameters

<i>Preg</i>	Specifies the structure to receive the compiled output of the <b>regcomp</b> subroutine.
<i>Pattern</i>	Contains the basic or extended regular expression to be compiled by the <b>regcomp</b> subroutine.  The default regular expression type for the <i>Pattern</i> parameter is a basic regular expression. An application can specify extended regular expressions with the <b>REG_EXTENDED</b> flag.
<i>CFlags</i>	Contains the bitwise inclusive OR of 0 or more flags for the <b>regcomp</b> subroutine. These flags are defined in the <b>regex.h</b> file:  <b>REG_EXTENDED</b> Uses extended regular expressions.  <b>REG_ICASE</b> Ignores case in match.  <b>REG_NOSUB</b> Reports only success or failure in the <b>regex</b> subroutine. If this flag is not set, the <b>regcomp</b> subroutine sets the <b>re_nsub</b> structure to the number of parenthetic expressions found in the <i>Pattern</i> parameter.  <b>REG_NEWLINE</b> Prohibits . (period) and nonmatching bracket expression from matching a new-line character. The ^ (circumflex) and \$ (dollar sign) will match the zero-length string immediately following or preceding a new-line character.

### Return Values

If successful, the **regcomp** subroutine returns a value of 0. Otherwise, it returns another value indicating the type of failure, and the content of the *Preg* parameter is undefined.

### Error Codes

The following macro names for error codes may be written to the **errno** global variable under error conditions:

**REG\_BADPAT** Indicates a basic or extended regular expression that is not valid.

<b>REG_ECOLLATE</b>	Indicates a collating element referenced that is not valid.
<b>REG_ETYPE</b>	Indicates a character class-type reference that is not valid.
<b>REG_EESCAPE</b>	Indicates a trailing \ in pattern.
<b>REG_ESUBREG</b>	Indicates a number in \digit is not valid or in error.
<b>REG_EBRACK</b>	Indicates a [] imbalance.
<b>REG_EPAREN</b>	Indicates a \(\) or () imbalance.
<b>REG_EBRACE</b>	Indicates a \{\} imbalance.
<b>REG_BADBR</b>	Indicates the content of \{\} is unusable: not a number, number too large, more than two numbers, or first number larger than second.
<b>REG_ERANGE</b>	Indicates an unusable end point in range expression.
<b>REG_ESPACE</b>	Indicates out of memory.
<b>REG_BADRPT</b>	Indicates a ? (question mark), * (asterisk), or + (plus sign) not preceded by valid basic or extended regular expression.

If the **regcomp** subroutine detects an illegal basic or extended regular expression, it can return either the **REG\_BADPAT** error code or another that more precisely describes the error.

## Examples

The following example illustrates how to match a string (specified in the *string* parameter) against an extended regular expression (specified in the *Pattern* parameter):

```
#include <sys/types.h>
#include <regex.h>
int
match(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);
}
```

In the preceding example, errors are treated as no match. When there is no match or error, the calling process can get details by calling the **regerror** subroutine.

## Related Information

The **regerror** (“regerror Subroutine”) subroutine, **regexec** (“regexec Subroutine” on page 35) subroutine, **regfree** (“regfree Subroutine” on page 38) subroutine.

Subroutines Overview and Understanding Internationalized Regular Expression Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## regerror Subroutine

### Purpose

Returns a string that describes the *ErrCode* parameter.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <regex.h>
```

```
size_t regerror (ErrCode, Preg, ErrBuf, ErrBuf_Size)
int ErrCode;
const regex_t * Preg;
char * ErrBuf;
size_t ErrBuf_Size;
```

## Description

The **regerror** subroutine provides a mapping from error codes returned by the **regcomp** and **regex** subroutines to printable strings. It generates a string corresponding to the value of the *ErrCode* parameter, which is the last nonzero value returned by the **regcomp** or **regex** subroutine with the given value of the *Preg* parameter. If the *ErrCode* parameter is not such a value, the content of the generated string is unspecified. The string generated is obtained from the **regex.cat** message catalog.

If the *ErrBuf\_Size* parameter is not 0, the **regerror** subroutine places the generated string into the buffer specifier by the *ErrBuf* parameter, whose size in bytes is specified by the *ErrBuf\_Size* parameter. If the string (including the terminating null character) cannot fit in the buffer, the **regerror** subroutine truncates the string and null terminates the result.

## Parameters

<i>ErrCode</i>	Specifies the error for which a description string is to be returned.
<i>Preg</i>	Specifies the structure that holds the previously compiled output of the <b>regcomp</b> subroutine.
<i>ErrBuf</i>	Specifies the buffer to receive the string generated by the <b>regerror</b> subroutine.
<i>ErrBuf_Size</i>	Specifies the size of the <i>ErrBuf</i> parameter.

## Return Values

The **regerror** subroutine returns the size of the buffer needed to hold the entire generated string, including the null termination. If the return value is greater than the value of the *ErrBuf\_Size* variable, the string returned in the *ErrBuf* buffer is truncated.

## Error Codes

If the *ErrBuf\_Size* value is 0, the **regerror** subroutine ignores the *ErrBuf* parameter, but returns the one of the following error codes. These error codes defined in the `regex.h` file.

<b>REG_NOMATCH</b>	Indicates the basic or extended regular expression was unable to find a match.
<b>REG_BADPAT</b>	Indicates a basic or extended regular expression that is not valid.
<b>REG_ECOLLATE</b>	Indicates a collating element referenced that is not valid.
<b>REG_ECTYPE</b>	Indicates a character class-type reference that is not valid.
<b>REG_EESCAPE</b>	Indicates a trailing <code>\</code> in pattern.
<b>REG_ESUBREG</b>	Indicates a number in <code>\digit</code> is not valid or in error.
<b>REG_EBRACK</b>	Indicates a <code>[]</code> imbalance.
<b>REG_EPAREN</b>	Indicates a <code>\(\)</code> or <code>()</code> imbalance.
<b>REG_EBRACE</b>	Indicates a <code>\{\}</code> imbalance.
<b>REG_BADBR</b>	Indicates the content of <code>\{\}</code> is unusable: not a number, number too large, more than two numbers, or first number larger than second.
<b>REG_ERANGE</b>	Indicates an unusable end point in range expression.
<b>REG_ESPACE</b>	Indicates out of memory.
<b>REG_BADRPT</b>	Indicates a <code>?</code> (question mark), <code>*</code> (asterisk), or <code>+</code> (plus sign) not preceded by valid basic or extended regular expression.

**REG\_ENEWLINE** Indicates a new-line character was found before the end of the regular or extended regular expression, and **REG\_NEWLINE** was not set.

If the *Preg* parameter passed to the **regexec** subroutine is not a compiled basic or extended regular expression returned by the **regcomp** subroutine, the result is undefined.

## Examples

An application can use the **regerror** subroutine (with the parameters (*Code*, *Preg*, null, (**size\_t**) 0) passed to it) to determine the size of buffer needed for the generated string, call the **malloc** subroutine to allocate a buffer to hold the string, and then call the **regerror** subroutine again to get the string. Alternately, this subroutine can allocate a fixed, static buffer that is large enough to hold most strings (perhaps 128 bytes), and then call the **malloc** subroutine to allocate a larger buffer if necessary.

## Related Information

The **regcomp** (“regcomp Subroutine” on page 32) subroutine, **regexec** (“regexec Subroutine”) subroutine, **regfree** (“regfree Subroutine” on page 38) subroutine.

Subroutines Overview and Understanding Internationalized Regular Expression Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## regexec Subroutine

### Purpose

Compares the null-terminated string specified by the value of the *String* parameter against the compiled basic or extended regular expression *Preg*, which must have previously been compiled by a call to the **regcomp** subroutine.

### Library

Standard C Library (**libc. a**)

### Syntax

```
#include <regex.h>

int regexec (Preg, String, NMatch, PMatch, EFlags)
const regex_t * Preg;
const char * String;
size_t NMatch;
regmatch_t * PMatch;
int EFlags;
```

### Description

The **regexec** subroutine compares the null-terminated string in the *String* parameter with the compiled basic or extended regular expression in the *Preg* parameter initialized by a previous call to the **regcomp** subroutine. If a match is found, the **regexec** subroutine returns a value of 0. The **regexec** subroutine returns a nonzero value if it finds no match or it finds an error.

If the *NMatch* parameter has a value of 0, or if the **REG\_NOSUB** flag was set on the call to the **regcomp** subroutine, the **regexec** subroutine ignores the *PMatch* parameter. Otherwise, the *PMatch* parameter points to an array of at least the number of elements specified by the *NMatch* parameter. The **regexec** subroutine fills in the elements of the array pointed to by the *PMatch* parameter with offsets of the

substrings of the *String* parameter. The offsets correspond to the parenthetic subexpressions of the original *pattern* parameter that was specified to the **regcomp** subroutine.

The **pmatch.rm\_so** structure is the byte offset of the beginning of the substring, and the **pmatch.rm\_eo** structure is one greater than the byte offset of the end of the substring. Subexpression *i* begins at the *i*th matched open parenthesis, counting from 1. The 0 element of the array corresponds to the entire pattern. Unused elements of the *PMatch* parameter, up to the value *PMatch*[*NMatch*-1], are filled with -1. If more than the number of subexpressions specified by the *NMatch* parameter (the *pattern* parameter itself counts as a subexpression), only the first *NMatch*-1 subexpressions are recorded.

When a basic or extended regular expression is being matched, any given parenthetic subexpression of the *pattern* parameter might match several different substrings of the *String* parameter. Otherwise, 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 the *PMatch* parameter when regular expressions are matched:

- If a subexpression in a regular expression participated in the match several times, the offset of the last matching substring is reported in the *PMatch* parameter.
- If a subexpression did not participate in a match, the byte offset in the *PMatch* parameter is a value of -1. A subexpression does not participate in a match if any of the following are true:
  - An \* (asterisk) or \{\} (backslash, left brace, backslash, right brace) appears immediately after the subexpression in a basic regular expression.
  - An \* (asterisk), ? (question mark), or { } (left and right braces) appears immediately after the subexpression in an extended regular expression and the subexpression did not match (matched 0 times).
  - A | (pipe) is used in an extended regular expression to select either the subexpression that didn't match or another subexpression, and the other subexpression matched.
- If a subexpression is contained in a subexpression, the data in the *PMatch* parameter refers to the last such subexpression.
- If a subexpression is contained in a subexpression and the byte offsets in the *PMatch* parameter have a value of -1, the pointers in the *PMatch* parameter also have a value of -1.
- If a subexpression matched a zero-length string, the offsets in the *PMatch* parameter refer to the byte immediately following the matching string.

If the **REG\_NOSUB** flag was set in the *cflags* parameter in the call to the **regcomp** subroutine, and the *NMatch* parameter is not equal to 0 in the call to the **regex** subroutine, the content of the *PMatch* array is unspecified.

If the **REG\_NEWLINE** flag was not set in the *cflags* parameter when the **regcomp** subroutine was called, then a new-line character in the *pattern* or *String* parameter is treated as an ordinary character. If the **REG\_NEWLINE** flag was set when the **regcomp** subroutine was called, the new-line character is treated as an ordinary character except as follows:

- A new-line character in the *String* parameter is not matched by a period outside of a bracket expression or by any form of a nonmatching list. A nonmatching list expression begins with a ^ (circumflex) and specifies a list that matches any character or collating element and the expression in the list after the leading caret. For example, the regular expression [^abc] matches any character except a, b, or c. The circumflex has this special meaning only when it is the first character in the list, immediately following the left bracket.
- A ^ (circumflex) in the *pattern* parameter, when used to specify expression anchoring, matches the zero-length string immediately after a new-line character in the *String* parameter, regardless of the setting of the **REG\_NOTBOL** flag.

- A \$ (dollar sign) in the *pattern* parameter, when used to specify expression anchoring, matches the zero-length string immediately before a new-line character in the *String* parameter, regardless of the setting of the **REG\_NOTEOL** flag.

## Parameters

<i>Preg</i>	Contains the compiled basic or extended regular expression to compare against the <i>String</i> parameter.
<i>String</i>	Contains the data to be matched.
<i>NMatch</i>	Contains the number of subexpressions to match.
<i>PMatch</i>	Contains the array of offsets into the <i>String</i> parameter that match the corresponding subexpression in the <i>Preg</i> parameter.
<i>EFlags</i>	Contains the bitwise inclusive OR of 0 or more of the flags controlling the behavior of the <b>regexec</b> subroutine capable of customizing.

The *EFlags* parameter modifies the interpretation of the contents of the *String* parameter. It is the bitwise inclusive OR of 0 or more of the following flags, which are defined in the **regex.h** file:

### **REG\_NOTBOL**

The first character of the string pointed to by the *String* parameter is not the beginning of the line. Therefore, the ^ (circumflex), when used as a special character, does not match the beginning of the *String* parameter.

### **REG\_NOTEOL**

The last character of the string pointed to by the *String* parameter is not the end of the line. Therefore, the \$ (dollar sign), when used as a special character, does not match the end of the *String* parameter.

## Return Values

On successful completion, the **regexec** subroutine returns a value of 0 to indicate that the contents of the *String* parameter matched the contents of the *pattern* parameter, or to indicate that no match occurred. The **REG\_NOMATCH** error is defined in the **regex.h** file.

## Error Codes

If the **regexec** subroutine is unsuccessful, it returns a nonzero value indicating the type of problem. The following macros for possible error codes that can be returned are defined in the **regex.h** file:

<b>REG_NOMATCH</b>	Indicates the basic or extended regular expression was unable to find a match.
<b>REG_BADPAT</b>	Indicates a basic or extended regular expression that is not valid.
<b>REG_ECOLLATE</b>	Indicates a collating element referenced that is not valid.
<b>REG_ECTYPE</b>	Indicates a character class-type reference that is not valid.
<b>REG_EESCAPE</b>	Indicates a trailing \ (backslash) in the pattern.
<b>REG_ESUBREG</b>	Indicates a number in \digit is not valid or is in error.
<b>REG_EBRACK</b>	Indicates a [ ] (left and right brackets) imbalance.
<b>REG_EPAREN</b>	Indicates a \ ( \ ) (backslash, left parenthesis, backslash, right parenthesis) or ( ) (left and right parentheses) imbalance.
<b>REG_EBRACE</b>	Indicates a \ { \ } (backslash, left brace, backslash, right brace) imbalance.
<b>REG_BADBR</b>	Indicates the content of \ { \ } (backslash, left brace, backslash, right brace) is unusable (not a number, number too large, more than two numbers, or first number larger than second).
<b>REG_ERANGE</b>	Indicates an unusable end point in range expression.
<b>REG_ESPACE</b>	Indicates out of memory.
<b>REG_BADRPT</b>	Indicates a ? (question mark), * (asterisk), or + (plus sign) not preceded by valid basic or extended regular expression.

If the value of the *Preg* parameter to the **regexec** subroutine is not a compiled basic or extended regular expression returned by the **regcomp** subroutine, the result is undefined.

## Examples

The following example demonstrates how the **REG\_NOTBOL** flag can be used with the **regexec** subroutine to find all substrings in a line that match a pattern supplied by a user. (For simplicity, very little error-checking is done in this example.)

```
(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.r_sp and pm.rm_ep>
/* This call to regexec finds the next match */
error = regexec (&re, pm.rm_ep, 1, &pm, REG_NOTBOL) ;
```

## Related Information

The **regcomp** (“regcomp Subroutine” on page 32) subroutine, **regerror** (“regerror Subroutine” on page 33) subroutine, **regfree** (“regfree Subroutine”) subroutine.

Subroutines Overview and Understanding Internationalized Regular Expression Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## regfree Subroutine

### Purpose

Frees any memory allocated by the **regcomp** subroutine associated with the *Preg* parameter.

### Library

Standard C Library (**libc. a**)

### Syntax

```
#include <regex.h>
```

```
void regfree ( Preg)
regex_t *Preg;
```

### Description

The **regfree** subroutine frees any memory allocated by the **regcomp** subroutine associated with the *Preg* parameter. An expression defined by the *Preg* parameter is no longer treated as a compiled basic or extended regular expression after it is given to the **regfree** subroutine.

### Parameters

*Preg* Structure containing the compiled output of the **regcomp** subroutine. Memory associated with this structure is freed by the **regfree** subroutine.

## Related Information

The **regcomp** (“regcomp Subroutine” on page 32) subroutine, **regerror** (“regerror Subroutine” on page 33) subroutine, **regexec** (“regexec Subroutine” on page 35) subroutine.

Subroutines Overview and Understanding Internationalized Regular Expression Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## reltimerid Subroutine

### Purpose

Releases a previously allocated interval timer.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int reltimerid ( TimerID)
timer_t TimerID;
```

### Description

The **reltimerid** subroutine is used to release a previously allocated interval timer, which is returned by the **gettimerid** subroutine. Any pending timer event generated by this interval timer is cancelled when the call returns.

### Parameters

*TimerID*        Specifies the ID of the interval timer being released.

### Return Values

The **reltimerid** subroutine returns a 0 if it is successful. If an error occurs, the value -1 is returned and **errno** is set.

### Error Codes

If the **reltimerid** subroutine fails, a -1 is returned and **errno** is set with the following error code:

**EINVAL**        The timer ID specified by the *Timerid* parameter is not a valid timer ID.

### Related Information

The **gettimerid** subroutine.

List of Time Data Manipulation Services in *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## remainder, remainderf, or remainderl Subroutine

### Purpose

Returns the floating-point remainder.

## Syntax

```
#include <math.h>
```

```
double remainder (x, y)
double x;
double y;
```

```
float remainderf (x, y)
float x;
float y;
```

```
long double remainderl (x, y)
long double x;
long double y ;
```

## Description

The **remainder**, **remainderf**, and **remainderl** subroutines return the floating-point remainder  $r=x - ny$  when  $y$  is nonzero. The value  $n$  is the integral value nearest the exact value  $x/y$ . When  $|n x/y| = \frac{1}{2}$ , the value  $n$  is chosen to be even.

## Parameters

$x$  Specifies the value of the numerator.  
 $y$  Specifies the value of the denominator.

## Return Values

Upon successful completion, the **remainder**, **remainderf**, and **remainderl** subroutines return the floating-point remainder  $r=x - ny$  when  $y$  is nonzero.

If  $x$  or  $y$  is NaN, a NaN is returned.

If  $x$  is infinite or  $y$  is 0 and the other is non-NaN, a domain error occurs, and a NaN is returned.

## Related Information

**abs** Subroutine, **feclearexcept** Subroutine, **fetestexcept** Subroutine, and **lldiv** Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

**math.h** in *AIX 5L Version 5.2 Files Reference*.

---

## remove Subroutine

### Purpose

Removes a file.

### Library

Standard C Library (**libc.a**)

## Syntax

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

```
int remove( FileName)
const char *FileName;
```

## Description

The **remove** subroutine makes a file named by *FileName* inaccessible by that name. An attempt to open that file using that name does not work unless you recreate it. If the file is open, the subroutine does not remove it.

If the file designated by the *FileName* parameter has multiple links, the link count of files linked to the removed file is reduced by 1.

## Parameters

*FileName*        Specifies the name of the file being removed.

## Return Values

Upon successful completion, the **remove** subroutine returns a value of 0; otherwise it returns a nonzero value.

## Related Information

The **link** subroutine, **rename** (“rename Subroutine” on page 42) subroutine.

The **link** or **unlink** (“unlink Subroutine” on page 423) command.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## remquo, remquof, or remquol Subroutine

### Purpose

Returns the floating-point remainder.

### Syntax

```
#include <math.h>
```

```
double remquo (x, y, quo)
double x;
double y;
int *quo;
```

```
float remquof (x, y, quo)
float x;
float y;
int *quo;
```

```
long double remquol (x, y, quo)
long double x;
long double y;
int *quo;
```

### Description

The **remquo**, **remquof**, and **remquol** subroutines compute the same remainder as the **remainder**, **remainderf**, and **remainderl** functions, respectively. In the object pointed to by *quo*, they store a value whose sign is the sign of  $x/y$  and whose magnitude is congruent modulo  $2^n$  to the magnitude of the integral quotient of  $x/y$ , where  $n$  is 3.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

<i>x</i>	Specifies the value of the numerator.
<i>y</i>	Specifies the value of the denominator.
<i>quo</i>	Points to the object where a value whose sign is the sign of <i>x/y</i> is stored.

## Return Values

The **remquo**, **remquof**, and **remquol** subroutines return *x* REM *y*.

If *x* or *y* is NaN, a NaN is returned.

If *x* is  $\pm\text{Inf}$  or *y* is zero and the other argument is non-NaN, a domain error occurs, and a NaN is returned.

## Related Information

“remainder, remainderf, or remainderl Subroutine” on page 39

**feclearexcept** Subroutine, **fetestexcept** Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

**math.h** in *AIX 5L Version 5.2 Files Reference*.

---

## rename Subroutine

### Purpose

Renames a directory or a file.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int rename ( FromPath, ToPath)  
const char *FromPath, *ToPath;
```

### Description

The **rename** subroutine renames a directory or a file within a file system.

To use the **rename** subroutine, the calling process must have write and search permission in the parent directories of both the *FromPath* and *ToPath* parameters. If the path defined in the *FromPath* parameter is a directory, the calling process must have write and search permission to the *FromPath* directory as well.

If the *FromPath* and *ToPath* parameters both refer to the same existing file, the **rename** subroutine returns successfully and performs no other action.

The components of both the *FromPath* and *ToPath* parameters must be of the same type (that is, both directories or both non-directories) and must reside on the same file system. If the *ToPath* file already exists, it is first removed. Removing it guarantees that a link named *ToPath* will exist throughout the operation. This link refers to the file named by either the *ToPath* or *FromPath* parameter before the operation began.

If the final component of the *FromPath* parameter is a symbolic link, the symbolic link (not the file or directory to which it points) is renamed. If the *ToPath* is a symbolic link, the link is destroyed.

If the parent directory of the *FromPath* parameter has the Sticky bit attribute (described in the **sys/mode.h** file), the calling process must have an effective user ID equal to the owner ID of the *FromPath* parameter, or to the owner ID of the parent directory of the *FromPath* parameter.

A user who is not the owner of the file or directory must have root user authority to use the **rename** subroutine.

If the *FromPath* and *ToPath* parameters name directories, the following must be true:

- The directory specified by the *FromPath* parameter is not an ancestor of *ToPath*. For example, the *FromPath* path name must not contain a path prefix that names the directory specified by the *ToPath* parameter.
- The directory specified in the *FromPath* parameter must be well-formed. A well-formed directory contains both . (dot) and .. (dot dot) entries. That is, the . (dot) entry in the *FromPath* directory refers to the same directory as that in the *FromPath* parameter. The .. (dot dot) entry in the *FromPath* directory refers to the directory that contains an entry for *FromPath*.
- The directory specified by the *ToPath* parameter, if it exists, must be well-formed (as defined previously).

## Parameters

<i>FromPath</i>	Identifies the file or directory to be renamed.
<i>ToPath</i>	Identifies the new path name of the file or directory to be renamed. If <i>ToPath</i> is an existing file or empty directory, it is replaced by <i>FromPath</i> . If <i>ToPath</i> specifies a directory that is not empty, the <b>rename</b> subroutine exits with an error.

## Return Values

Upon successful completion, the **rename** subroutine returns a value of 0. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **rename** subroutine is unsuccessful and the file or directory name remains unchanged if one or more of the following are true:

<b>EACCES</b>	Creating the requested link requires writing in a directory mode that denies the process write permission.
<b>EBUSY</b>	The directory named by the <i>FromPath</i> or <i>ToPath</i> parameter is currently in use by the system, or the file named by <i>FromPath</i> or <i>ToPath</i> is a named STREAM.
<b>EDQUOT</b>	The directory that would contain the path specified by the <i>ToPath</i> parameter cannot be extended because the user's or group's quota of disk blocks on the file system containing the directory is exhausted.
<b>EEXIST</b>	The <i>ToPath</i> parameter specifies an existing directory that is not empty.
<b>EINVAL</b>	The path specified in the <i>FromPath</i> or <i>ToPath</i> parameter is not a well-formed directory ( <i>FromPath</i> is an ancestor of <i>ToPath</i> ), or an attempt has been made to rename . (dot) or .. (dot dot).

<b>EISDIR</b>	The <i>ToPath</i> parameter names a directory and the <i>FromPath</i> parameter names a non-directory.
<b>EMLINK</b>	The <i>FromPath</i> parameter names a directory that is larger than the maximum link count of the parent directory of the <i>ToPath</i> parameter.
<b>ENOENT</b>	A component of either path does not exist, the file named by the <i>FromPath</i> parameter does not exist, or a symbolic link was named, but the file to which it refers does not exist.
<b>ENOSPC</b>	The directory that would contain the path specified in the <i>ToPath</i> parameter cannot be extended because the file system is out of space.
<b>ENOTDIR</b>	The <i>FromPath</i> parameter names a directory and the <i>ToPath</i> parameter names a non-directory.
<b>ENOTEMPTY</b>	The <i>ToPath</i> parameter specifies an existing directory that is not empty.
<b>EROFS</b>	The requested operation requires writing in a directory on a read-only file system.
<b>ETXTBSY</b>	The <i>ToPath</i> parameter names a shared text file that is currently being used.
<b>EXDEV</b>	The link named by the <i>ToPath</i> parameter and the file named by the <i>FromPath</i> parameter are on different file systems.

If Network File System (NFS) is installed on the system, the **rename** subroutine can be unsuccessful if the following is true:

**ETIMEDOUT**      The connection timed out.

The **rename** subroutine can be unsuccessful for other reasons. See Appendix A, "Base Operating System Error Codes For Services That Require Path-Name Resolution" for a list of additional errors.

## Related Information

The **chmod** subroutine, **link** subroutine, **mkdir** subroutine, **rmdir** ("rmdir Subroutine" on page 48) subroutine, **unlink** ("unlink Subroutine" on page 423) subroutine.

The **chmod** command, **mkdir** command, **mv** command, **mmdir** command.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

## reset\_malloc\_log Subroutine

### Purpose

Resets information collected by the malloc subsystem.

### Syntax

```
#include <malloc.h>
void reset_malloc_log (addr)
void *addr;
```

### Description

The **reset\_malloc\_log** subroutine resets the record of currently active malloc allocations stored by the malloc subsystem. These records are stored in **malloc\_log** structures, which are located in the process heap. Only records corresponding to the heap of which *addr* is a member are reset, unless *addr* is NULL, in which case records for all heaps are reset. The *addr* parameter must be a pointer to space allocated previously by the malloc subsystem or NULL, otherwise no information is reset and the **errno** global variable is set to **EINVAL**.

## Parameters

*addr* Pointer to space allocated previously by the malloc subsystem

## Related Information

malloc Subroutine, get\_malloc\_log Subroutine, and get\_malloc\_log\_live Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*

---

## revoke Subroutine

### Purpose

Revokes access to a file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int revoke ( Path )
char *Path;
```

### Description

The **revoke** subroutine revokes access to a file by all processes.

All accesses to the file are revoked. Subsequent attempts to access the file using a file descriptor established before the **revoke** subroutine fail and cause the process to receive a return value of -1, and the **errno** global variable is set to **EBADF**.

A process can revoke access to a file only if its effective user ID is the same as the file owner ID, or if the calling process is privileged.

**Note:** The **revoke** subroutine has no effect on subsequent attempts to open the file. To assure exclusive access to the file, the caller should change the access mode of the file before issuing the **revoke** subroutine. Currently the **revoke** subroutine works only on terminal devices. The **chmod** subroutine changes file access modes.

## Parameters

*Path* Path name of the file for which access is to be revoked.

## Return Values

Upon successful completion, the **revoke** subroutine returns a value of 0.

If the **revoke** subroutine fails, a value of -1 returns and the **errno** global variable is set to indicate the error.

## Error Codes

The **revoke** subroutine fails if any of the following are true:

<b>ENOTDIR</b>	A component of the path prefix is not a directory.
<b>EACCES</b>	Search permission is denied on a component of the path prefix.

<b>ENOENT</b>	A component of the path prefix does not exist, or the process has the <b>disallow truncation</b> attribute (see the <b>ulimit</b> subroutine).
<b>ENOENT</b>	The path name is null.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ESTALE</b>	The process's root or current directory is located in a virtual file system that has been unmounted.
<b>EFAULT</b>	The <i>Path</i> parameter points outside of the process's address space.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the path name.
<b>ENAMETOOLONG</b>	A component of a path name exceeds 255 characters, or an entire path name exceeds 1023 characters.
<b>EIO</b>	An I/O error occurred during the operation.
<b>EPERM</b>	The effective user ID of the calling process is not the same as the file's owner ID.
<b>EINVAL</b>	Access rights revocation is not implemented for this file.

## Related Information

The **chmod** subroutine, **frevoke** subroutine.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## rintf, rintl, or rint Subroutine

### Purpose

Rounds to the nearest integral value.

### Syntax

```
#include <math.h>
```

```
float rintf (x)
float x;
```

```
long double rintl (x)
long double x;
```

```
double rint (x)
double x;
```

### Description

The **rintf**, **rintl**, and **rint** subroutines return the integral value (represented as a double) nearest *x* in the direction of the current rounding mode. The current rounding mode is implementation-defined.

The **rintf**, **rintl**, and **rint** subroutines differ from the **nearbyint**, **nearbyintf**, and **nearbyintl** subroutines only in that they may raise the inexact floating-point exception if the result differs in value from the argument.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

*x* Specifies the value to be rounded.

## Return Values

Upon successful completion, the **rintf**, **rintl**, and **rint** subroutines return the integer (represented as a double precision number) nearest  $x$  in the direction of the current rounding mode.

If  $x$  is NaN, a NaN is returned.

If  $x$  is  $\pm 0$  or  $\pm \text{Inf}$ ,  $x$  is returned.

If the correct value would cause overflow, a range error occurs the **rintf**, **rintl**, and **rint** subroutines return the value of the macro  **$\pm\text{HUGE\_VALF}$**  and  **$\pm\text{HUGE\_VALL}$**  (with the same sign as  $x$ ), respectively.

## Related Information

**abs** Subroutine, **floor**, **floorl**, **ceil**, **ceill**, **nearest**, **trunc**, **rint**, **itrunc**, **uitrunc**, **fmod**, **fmodl**, **fabs**, or **fabsl** Subroutine, **feclearexcept** Subroutine, **fetestexcept** Subroutine, **class**, **\_class**, **finite**, **isnan**, or **unordered** Subroutines, and **lldiv** Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

**math.h** in *AIX 5L Version 5.2 Files Reference*.

---

## round, roundf, or roundl Subroutine

### Purpose

Rounds to the nearest integer value in a floating-point format.

### Syntax

```
#include <math.h>
```

```
double round (x)
double x;
```

```
float roundf (x)
float x;
```

```
long double roundl (x)
long double x;
```

### Description

The **round**, **roundf**, and **roundl** subroutines round the  $x$  parameter to the nearest integer value in floating-point format, rounding halfway cases away from zero, regardless of the current rounding direction.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

$x$  Specifies the value to be rounded.

## Return Values

Upon successful completion, the **round**, **roundf**, and **roundl** subroutines return the rounded integer value.

If  $x$  is NaN, a NaN is returned.

If  $x$  is  $\pm 0$  or  $\pm \text{Inf}$ ,  $x$  is returned.

If the correct value would cause overflow, a range error occurs and the **round**, **roundf**, and **roundl** subroutines return the value of the macro  **$\pm\text{HUGE\_VAL}$** ,  **$\pm\text{HUGE\_VALF}$** , and  **$\pm\text{HUGE\_VALL}$**  (with the same sign as  $x$ ), respectively.

## Related Information

feclearexcept Subroutine and fetestexcept Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

**math.h** in *AIX 5L Version 5.2 Files Reference*.

---

## rmdir Subroutine

### Purpose

Removes a directory.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int rmdir ( Path )  
const char *Path;
```

### Description

The **rmdir** subroutine removes the directory specified by the *Path* parameter. If Network File System (NFS) is installed on your system, this path can cross into another node.

For the **rmdir** subroutine to execute successfully, the calling process must have write access to the parent directory of the *Path* parameter.

In addition, if the parent directory of *Path* has the Sticky bit attribute (described in the **sys/mode.h** file), the calling process must have one of the following:

- An effective user ID equal to the directory to be removed
- An effective user ID equal to the owner ID of the parent directory of *Path*
- Root user authority.

## Parameters

*Path* Specifies the directory path name. The directory you specify must be:

**Empty** The directory contains no entries other than . (dot) and .. (dot dot).

**Well-formed**

If the . (dot) entry in the *Path* parameter exists, it must refer to the same directory as *Path*.

Exactly one directory has a link to the *Path* parameter, excluding the self-referential . (dot). If the .. (dot dot) entry in *Path* exists, it must refer to the directory that contains an entry for *Path*.

## Return Values

Upon successful completion, the **rmdir** subroutine returns a value of 0. Otherwise, a value of -1 is returned, the specified directory is not changed, and the **errno** global variable is set to indicate the error.

## Error Codes

The **rmdir** subroutine fails and the directory is not deleted if the following errors occur:

<b>EACCES</b>	There is no search permission on a component of the path prefix, or there is no write permission on the parent directory of the directory to be removed.
<b>EBUSY</b>	The directory is in use as a mount point.
<b>EEXIST</b> or <b>ENOTEMPTY</b>	The directory named by the <i>Path</i> parameter is not empty.
<b>ENAMETOOLONG</b>	The length of the <i>Path</i> parameter exceeds <b>PATH_MAX</b> ; or a path-name component longer than <b>NAME_MAX</b> and <b>POSIX_NO_TRUNC</b> is in effect.
<b>ENOENT</b>	The directory named by the <i>Path</i> parameter does not exist, or the <i>Path</i> parameter points to an empty string.
<b>ENOTDIR</b>	A component specified by the <i>Path</i> parameter is not a directory.
<b>EINVAL</b>	The directory named by the <i>Path</i> parameter is not well-formed.
<b>EROFS</b>	The directory named by the <i>Path</i> parameter resides on a read-only file system.

The **rmdir** subroutine can be unsuccessful for other reasons. See Appendix A, "Base Operating System Error Codes For Services That Require Path-Name Resolution" on page A-1 for a list of additional errors.

If NFS is installed on the system, the **rmdir** subroutine fails if the following is true:

**ETIMEDOUT** The connection timed out.

## Related Information

The **chmod** or **fchmod** subroutine, **mkdir** subroutine, **remove** ("remove Subroutine" on page 40) subroutine, **rename** ("rename Subroutine" on page 42) subroutine, **umask** ("umask Subroutine" on page 418) subroutine, **unlink** ("unlink Subroutine" on page 423) subroutine.

The **rm** command, **rmdir** command.

Files, Directories, and File Systems For Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## rpmatch Subroutine

### Purpose

Determines whether the response to a question is affirmative or negative.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <stdlib.h>

int rpmatch ( Response)
const char *Response;
```

## Description

The **rpmatch** subroutine determines whether the expression in the *Response* parameter matches the affirmative or negative response specified by the **LC\_MESSAGES** category in the current locale. Both expressions can be extended regular expressions.

## Parameters

*Response*        Specifies input entered in response to a question that requires an affirmative or negative reply.

## Return Values

This subroutine returns a value of 1 if the expression in the *Response* parameter matches the locale's affirmative expression. It returns a value of 0 if the expression in the *Response* parameter matches the locale's negative expression. If neither expression matches the expression in the *Response* parameter, a -1 is returned.

## Examples

The following example shows an affirmative expression in the *En\_US* locale. This example matches any expression in the *Response* parameter that begins with a y or Y followed by zero or more alphabetic characters, or it matches the letter o followed by the letter k.

```
^[yY][:alpha:]* | ok
```

## Related Information

The **localeconv** subroutine, **nl\_langinfo** subroutine, **regcomp** (“regcomp Subroutine” on page 32) subroutine, **regex** (“regex Subroutine” on page 35) subroutine, **setlocale** (“setlocale Subroutine” on page 136) subroutine.

National Language Support Overview and Setting the Locale in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## RSiAddSetHot Subroutine

### Purpose

Add a single set of peer statistics to an already defined **SpmiHotSet**.

### Library

RSI Library (**libSpmi.a**)

## Syntax

```
#include sys/Rsi.h

struct SpmiHotVals *RSiAddSetHot(rhandle, HotSet, StatName,
GrandParent,
                                maxresp, threshold, frequency, feed_type,
                                except_type, severity, trap_no)

RSiHandle rhandle;
struct SpmiHotSet *HotSet;
char *StatName;
cx_handle GrandParent;
int maxresp;
int threshold;
int frequency;
int feed_type;
int excp_type;
int severity;
int trap_no;
```

## Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**HotSet** Specifies a pointer to a valid structure of type **SpmiHotSet** as created by the **RSiCreateHotSet** (“**RSiCreateHotSet Subroutine**” on page 56) subroutine call.

**StatName** Specifies the name of the statistic within the subcontexts (peer contexts) of the context identified by the *GrandParent* parameter.

**GrandParent** Specifies a valid **cx\_handle** handle as obtained by another subroutine call. The handle must identify a context with at least one subcontext, which contains the statistic identified by the *StatName* parameter. If the context specified is one of the **RTime** contexts, no subcontext need to be created at the time the **SpmiAddSetHot** subroutine call is issued; the presence of the metric identified by the *StatName* parameter is checked against the context class description.

If the context specified has or may have multiple levels of instantiable context below it (such as the **FS** and **RTime/ARM** contexts), the metric is only searched for at the lowest context level. The **SpmiHotSet** created is a pseudo hotvals structure used to link together a peer group of **SpmiHotVals** structures, which are created under the covers, one for each subcontext of the *GrandParent* context. In the case of **RTime/ARM**, if additional contexts are later added under the *GrandParent* contexts, additional hotsets are added to the peer group. This is transparent to the application program, except that the **RSiGetHotItem** (“**RSiGetHotItem Subroutine**” on page 63) subroutine call will return the peer group **SpmiHotVals** pointer rather than the pointer to the pseudo structure.

Note that specifying a specific volume group context (such as **FS/rootvg**) or a specific application context (such as **RTime/ARN/armpeek**) is still valid and won't involve creation of pseudo **SpmiHotVals** structures.

**maxresp** Must be non-zero if *excp\_type* specifies that exceptions or SNMP traps must be generated. If specified as zero, indicates that all **SpmiHotItems** that meet the criteria specified by *threshold* must be returned, up-to a maximum of *maxresp* items. If both exceptions/traps and feeds are requested, the *maxresp* value is used to cap the number of exceptions/alerts as well as the number of items returned. If *feed\_type* is specified as **SiHotAlways**, the *maxresp* parameter is still used to return at most *maxresp* items.

Where the *GrandParent* argument specifies a context that has multiple levels of instantiable contexts below it, the *maxresp* is applied to each of the lowest level contexts above the the actual peer contexts at a time. For example, if the *GrandParent* context is **FS** (file systems) and the system has three volume groups, then a *maxresp* value of 2 could cause up to a maximum of 2 x 3 = 6 responses to be generated.

**threshold** Must be non-zero if *excp\_type* specifies that exceptions or SNMP traps must be generated. If specified as zero, indicates that all values read qualify to be returned in feeds. The value specified is compared to the data value read for each peer statistic. If the data value exceeds the *threshold*, it qualifies to be returned as an **SpmiHotItems** element in the **SpmiHotVals** structure. If the *threshold* is specified as a negative value, the value qualifies if it is lower than the numeric value of *threshold*. If *feed\_type* is specified as **SiHotAlways**, the threshold value is ignored for feeds. For peer statistics of type **SiCounter**, the *threshold* must be specified as a rate per second; for **SiQuantity** statistics the *threshold* is specified as a level.

**frequency** Must be non-zero if *excp\_type* specifies that exceptions or SNMP traps must be generated. Ignored for feeds. Specifies the minimum number of minutes that must expire between any two exceptions/traps generated from this **SpmiHotVals** structure. This value must be specified as no less than 5 minutes.

**feed\_type** Specifies if feeds of **SpmiHotItems** should be returned for this **SpmiHotVals** structure. The following values are valid:

- **SiHotNoFeed** No feeds should be generated
- **SiHotThreshold** Feeds are controlled by *threshold*.
- **SiHotAlways** All values, up-to a maximum of *maxresp* must be returned as feeds.

**excp\_type** Controls the generation of exception data packets and/or the generation of SNMP Traps from **xmservd**. Note that these types of packets and traps can only actually be sent if **xmservd** is running. Because of this, exception packets and SNMP traps are only generated as long as **xmservd** is active. Traps can only be generated on AIX. The conditions for generating exceptions and traps are controlled by the *threshold* and *frequency* parameters. The following values are valid for *excp\_type*:

- **SiNoHotException** Generate neither exceptions nor traps.
- **SiHotException** Generate exceptions but not traps.
- **SiHotTrap** Generate SNMP traps but not exceptions.
- **SiHotBoth** Generate both exceptions and SNMP traps.

**severity** Required to be positive and greater than zero if exceptions are generated, otherwise specify as zero. Used to assign a severity code to the exception for display by **exmon**.

**trap\_no** Required to be positive and greater than zero if SNMP traps are generated, otherwise specify as zero. Used to assign the trap number in the generated SNMP trap.

This subroutine is part of the Performance Toolbox for AIX licensed product.

## Return Values

If successful, the subroutine returns a pointer to a structure of type **struct SpmiHotVals**. If an error occurs, NULL is returned and an error text may be placed in the external character array **RSiEMsg**. If you attempt to add more values to a statset than the current local buffer size allows, **RSiErrno** is set to **RSiTooMany**. If you attempt to add more values than the buffer size of the remote host's **xmservd** daemon allows, **RSiErrno** is set to **RSiBadStat** and the status field in the returned packet is set to **too\_many\_values**.

The external integer **RSiMaxValues** holds the maximum number of values acceptable with the data-consumer's buffer size.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];

- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSi error codes are described in *List of RSi Error Codes* .

## Files

<code>/usr/include/sys/Rsi.h</code>	Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.
-------------------------------------	---

## Related Information

For related information, see:

- “RSiCreateHotSet Subroutine” on page 56
- “RSiOpen Subroutine” on page 75.

## RSiChangeFeed Subroutine

### Purpose

Changes the frequency at which the **xmservd** on the host identified by the first argument daemon is sending **data\_feed** packets for a statset.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
int RSiChangeFeed(rhandle, statset, msec)
RSiHandle rhandle; struct SpmiStatSet *statset; int msec;
```

### Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**statset** Must be a pointer to a structure of type **struct SpmiStatSet**, which was previously returned by a successful **RSiCreateStatSet** subroutine call. Data feeding must have been started for this **SpmiStatSet** via a previous **RSiStartFeed** (“**RSiStartFeed Subroutine**” on page 80) subroutine call.

**msec** The number of milliseconds between the sending of **data\_feed** packets. This number is rounded to a multiple of **min\_remote\_int** milliseconds by the **xmservd** daemon on the remote host. This minimum interval can be modified through the **-i** command line interval to **xmservd**.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Return Values

If successful, the subroutine returns zero, otherwise -1. A NULL error text is placed in the external character array **RSiEMsg** regardless of the subroutine’s success or failure.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- `extern char RSiEMsg[];`
- `extern int RSiErrno;`

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSI Error Codes*.

## Files

`/usr/include/sys/Rsi.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “RSiCreateStatSet Subroutine” on page 57
- “RSiOpen Subroutine” on page 75
- “RSiStartFeed Subroutine” on page 80.

---

## RSiChangeHotFeed Subroutine

### Purpose

Changes the frequency at which the **xmservd** on the host identified by the first argument daemon is sending **hot\_feed** packets for a statset or checking if exceptions or SNMP traps should be generated.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
int RSiChangeFeed(rhandle, hotset, msecs)
RSiHandle rhandle; struct SpmiHotSet *hotset; int msecs;
```

### Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**hotset** Must be a pointer to a structure of type **struct SpmiHotSet**, which was previously returned by a successful **RSiCreateHotSet** (“**RSiCreateHotSet Subroutine**” on page 56) subroutine call. Data feeding must have been started for this **SpmiHotSet** via a previous **RSiStartHotFeed** (“**RSiStartHotFeed Subroutine**” on page 81) subroutine call.

**msecs** The number of milliseconds between the sending of **Hot\_feed** packets. This number is rounded to a multiple of **min\_remote\_int** milliseconds by the **xmservd** daemon on the remote host. This minimum interval can be modified through the **-i** command line interval to **xmservd**.

This subroutine is part of the Performance Toolbox for AIX licensed product.

## Return Values

If successful, the subroutine returns zero, otherwise -1. A NULL error text is placed in the external character array **RSiEMsg** regardless of the subroutine's success or failure.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSi error codes are described in *List of RSi Error Codes*.

## Files

**/usr/include/sys/Rsi.h** Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

In the sample program, the **SpmiStatSet** is created in the local function **lststats** shown previously in lines 6 through 10.

- “RSiCreateHotSet Subroutine” on page 56
- “RSiOpen Subroutine” on page 75
- “RSiStartHotFeed Subroutine” on page 81.

---

## RSiClose Subroutine

### Purpose

Terminates the RSI interface for a remote host connection.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
void RSiClose(rhandle)
RSiHandle rhandle;
```

### Description

The **RSiClose** subroutine is responsible for:

1. Removing the data-consumer program as a known data consumer on a particular host. This is done by sending a **going\_down** packet to the host.
2. Marking the RSI handle as not active.
3. Releasing all memory allocated in connection with the RSI handle.
4. Terminating the RSI interface for a remote host.

A successful **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine creates tables on the remote host it was issued against. Therefore, a data consumer program that has issued successful **RSiOpen**

subroutine calls should issue an **RSiClose** (“**RSiClose Subroutine**” on page 55) subroutine call for each **RSiOpen** call before the program exits so that the tables in the remote **xmservd** daemon can be released.

This subroutine is part of the Performance Toolbox for AIX licensed product.

## Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** subroutine.

The macro **RSiIsOpen** can be used to test whether an RSI handle is open. It takes an **RSiHandle** as argument and returns true (1) if the handle is open, otherwise false (0).

## Files

**/usr/include/sys/Rsi.h** Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “RSiInit Subroutine” on page 68
- “RSiOpen Subroutine” on page 75

---

## RSiCreateHotSet Subroutine

### Purpose

Creates an empty hotset on the remote host identified by the argument.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
struct SpmiHotSet *RSiCreateHotSet(rhandle)
RSiHandle rhandle;
```

### Description

The **RSiCreateHotSet** subroutine allocates an **SpmiHotSet** structure. The structure is initialized as an empty **SpmiHotSet** and a pointer to the **SpmiHotSet** structure is returned.

The **SpmiHotSet** structure provides the anchor point to a set of peer statistics and must exist before the **RSiAddSetHot** (“**RSiAddSetHot Subroutine**” on page 50) subroutine can be successfully called.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

### Return Values

The **RSiCreateHotSet** subroutine returns a pointer to a structure of type **SpmiHotSet** if successful. If unsuccessful, the subroutine returns a NULL value.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- `extern char RSiEMsg[];`
- `extern int RSiErrno;`

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSi Error Codes* .

## Files

`/usr/include/sys/Rsi.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the **RSI**.

## Related Information

For related information, see:

- “RSiOpen Subroutine” on page 75
- “RSiAddSetHot Subroutine” on page 50.

---

## RSiCreateStatSet Subroutine

### Purpose

Creates an empty statset on the remote host identified by the argument.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
struct SpmiStatSet *RSiCreateStatSet(rhandle)
RSiHandle rhandle;
```

### Description

The **RSiCreateStatSet** subroutine allocates an **SpmiStatSet** structure. The structure is initialized as an empty **SpmiStatSet** and a pointer to the **SpmiStatSet** structure is returned.

The **SpmiStatSet** structure provides the anchor point to a set of statistics and must exist before the **RSiPathAddSetStat** (“**RSiPathAddSetStat Subroutine**” on page 78) subroutine can be successfully called.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

## Return Values

The **RSiCreateStatSet** subroutine returns a pointer to a structure of type **SpmiStatSet** if successful. If unsuccessful, the subroutine returns a NULL value.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSi error codes are described in *List of RSi Error Codes*.

## Files

**/usr/include/sys/Rsi.h** Declares the subroutines, data structures, handles, and macros that an application program can use to access the **RSI**.

## Related Information

For related information, see:

- “RSiOpen Subroutine” on page 75
- “RSiPathAddSetStat Subroutine” on page 78.

---

## RSiDelSetHot Subroutine

### Purpose

Deletes a single set of peer statistics identified by an **SpmiHotVals** structure from an **SpmiHotSet**.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
int RSiDelSetHot(rhandle, hsp, hvp)
RSiHandle rhandle; struct SpmiHotSet *hsp; struct SpmiHotVals *hvp;
```

### Description

The **RSiDelSetHot** subroutine performs the following actions:

1. Validates that the **SpmiHotSet** identified by the second argument exists and contains the **SpmiHotVals** statistic identified by the third argument.
2. Deletes the **SpmiHotVals** value from the **SpmiHotSet** so that future **data\_feed** packets do not include the deleted statistic.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**hsp** Must be a pointer to a structure type **struct SpmiHotSet**, which was previously returned by a successful **RSiCreateHotSet** subroutine call.

**hvp** Must be a handle of type **struct SpmiHotVals** as returned by a successful **RSiAddSetHot** (“**RSiAddSetHot Subroutine**” on page 50) subroutine call. You cannot specify an **SpmiHotVals** that was internally generated by the Spmi library code as described under the *GrandParent* parameter to **RSiAddSetHot** (“**RSiAddSetHot Subroutine**” on page 50).

## Return Values

If successful, the subroutine returns a zero value; otherwise it returns a non-zero value and an error text may be placed in the external character array **RSiErrMsg**.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiErrMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiErrMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSi error codes are described in *List of RSi Error Codes*.

## Files

`/usr/include/sys/Rsi.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “RSiOpen Subroutine” on page 75
- “RSiAddSetHot Subroutine” on page 50.

---

## RSiDelSetStat Subroutine

### Purpose

Deletes a single statistic identified by an **SpmiStatVals** pointer from an **SpmiStatSet**.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
int RSiDelSetStat(rhandle, ssp, svp)
RSiHandle rhandle; struct SpmiStatSet *ssp; struct SpmiStatVals *svp;
```

### Description

The **RSiDelSetStat** subroutine performs the following actions:

1. Validates the **SpmiStatSet** identified by the second argument exists and contains the **SpmiStatVals** statistic identified by the third argument.

2. Deletes the **SpmiStatVals** value from the **SpmiStatSet** so that future **data\_feed** packets do not include the deleted statistic.

This subroutine is part of the Performance Toolbox for AIX licensed product.

## Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**ssp** Must be a pointer to a structure type **struct SpmiStatSet**, which was previously returned by a successful **RSiCreateStatSet** (“**RSiCreateStatSet Subroutine**” on page 57) subroutine call.

**svp** Must be a handle of type **struct SpmiStatVals** as returned by a successful **RSiPathAddSetStat** (“**RSiPathAddSetStat Subroutine**” on page 78) subroutine call.

## Return Values

If successful, the subroutine returns a zero value; otherwise it returns a non-zero value and an error text may be placed in the external character array **RSiErrMsg**.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiErrMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiErrMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSI Error Codes*.

## Files

**/usr/include/sys/Rsi.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “**RSiCreateStatSet Subroutine**” on page 57
- “**RSiOpen Subroutine**” on page 75
- “**RSiPathAddSetStat Subroutine**” on page 78.

---

## RSiFirstCx Subroutine

### Purpose

Returns the first subcontext of an **SpmiCx** context.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
```

```

struct SpmiCxLink *RSiFirstCx(rhandle, context, name,
descr)
RSiHandle rhandle;
cx_handle *context;
char **name;
char **descr;

```

## Description

The **RSiFirstCx** subroutine performs the following actions:

1. Validates that the context identified by the second argument exists.
2. Returns a handle to the first element of the list of subcontexts defined for the context.
3. Returns the short name and description of the subcontext.

This subroutine is part of the Performance Toolbox for AIX licensed product.

## Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**context** Must be a handle of type **cx\_handle**, which was previously returned by a successful **RSiPathGetCx** (“**RSiPathGetCx Subroutine**” on page 79) subroutine call.

**name** Must be a pointer to a pointer to a character array. The pointer must be initialized to point at a character array pointer. When the subroutine call is successful, the short name of the subcontext is returned in the character array pointer.

**descr** Must be a pointer to a pointer to a character array. The pointer must be initialized to point at a character array pointer. When the subroutine call is successful, the description of the subcontext is returned in the character array pointer.

## Return Values

If successful, the subroutine returns a pointer to a structure of type **struct SpmiCxLink**. If an error occurs or if the context doesn't contain subcontexts, NULL is returned and an error text may be placed in the external character array **RSiErrMsg**.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiErrMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiErrMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSI Error Codes*.

## Files

**/usr/include/sys/Rsi.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “RSiNextCx Subroutine” on page 73
- “RSiOpen Subroutine” on page 75
- “RSiPathGetCx Subroutine” on page 79.

---

## RSiFirstStat Subroutine

### Purpose

Returns the first statistic of an **SpmiCx** context.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h

struct SpmiStatLink *RSiFirstStat(rhandle, context, name,
descr)
RSiHandle rhandle;
cx_handle context;
char **name;
char **descr;
```

### Description

The **RSiFirstStat** subroutine performs the following actions:

1. Validates that the context identified by the second argument exists.
2. Returns a handle to the first element of the list of statistics defined for the context.
3. Returns the short name and description of the statistic.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**context** Must be a handle of type **cx\_handle**, which was previously returned by a successful **RSiPathGetCx** (“**RSiPathGetCx Subroutine**” on page 79) subroutine call.

**name** Must be a pointer to a pointer to a character array. The pointer must be initialized to point at a character array pointer. When the subroutine call is successful, the short name of the statistics value is returned in the character array pointer.

**descr** Must be a pointer to a pointer to a character array. The pointer must be initialized to point at a character array pointer. When the subroutine call is successful, the description of the statistics value is returned in the character array pointer.

### Return Values

If successful, the subroutine returns a pointer to a structure of type **struct SpmiStatLink**. If an error occurs, NULL is returned and an error text may be placed in the external character array **RSiEMsg**.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- `extern char RSiEMsg[];`
- `extern int RSiErrno;`

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSi Error Codes*.

## Files

`/usr/include/sys/Rsi.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “RSiNextStat Subroutine” on page 74
- “RSiOpen Subroutine” on page 75
- “RSiPathGetCx Subroutine” on page 79.

---

## RSiGetHotItem Subroutine

### Purpose

Locates and decodes the next **SpmiHotItems** element at the current position in an incoming data packet of type **hot\_feed**.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
struct SpmiHotVals *RSiGetHotItem(rhandle, HotSet, index, value,
absvalue, name)
RSiHandle rhandle;
struct SpmiHotSet **HotSet;
int *index;
float *value;
float absvalue;
char **name;
```

### Description

The **RSiGetHotItem** subroutine locates the **SpmiHotItems** structure in the **hot\_feed** data packet indexed by the value of the *index* parameter. The subroutine returns a NULL value if no further **SpmiHotItems** structures are found. The **RSiGetHotItem** subroutine should only be executed after a successful call to the **RSiGetHotSet** subroutine.

The **RSiGetHotItem** subroutine is designed to be used for walking all **SpmiHotItems** elements returned in a **hot\_feed** data packet. Because the data packet may contain elements belonging to more than one **SpmiHotSet**, the *index* is purely abstract and is only used to keep position. By feeding the updated integer pointed to by *index* back to the next call, the walking of the **hot\_feed** packet can be done in a tight loop.

Successful calls to **RSiGetHotItem** will decode each **SpmiHotItems** element and return the data value in *value* and the name of the peer context that owns the corresponding statistic in *name*.

This subroutine is part of the Performance Toolbox for AIX licensed product.

## Parameters

<b>rhandle</b>	Must be an <b>RSiHandle</b> , which was previously initialized by the <b>RSiOpen</b> (“ <b>RSiOpen Subroutine</b> ” on page 75) subroutine.
<b>HotSet</b>	Used to return a pointer to a valid <b>SpmiHotSet</b> structure as obtained by a previous <b>RSiCreateHotSet</b> (“ <b>RSiCreateHotSet Subroutine</b> ” on page 56) subroutine call. The calling program can use this value to locate the <b>SpmiHotSet</b> if its address was stored by the program after it was created. The time stamps in the <b>SpmiHotSet</b> are updated with the time stamps of the decoded <b>SpmiHotItems</b> element.
<b>index</b>	A pointer to an integer that contains the desired relative element number in the <b>SpmiHotItems</b> array across all <b>SpmiStatVals</b> contained in the data packet. A value of zero points to the first element. When the <b>RSiGetHotItem</b> subroutine returns, the integer contain the index of the next <b>SpmiHotItems</b> element in the data packet. By passing the returned <i>index</i> parameter to the next call to <b>RSiGetHotItem</b> , the calling program can iterate through all <b>SpmiHotItems</b> elements in the <b>hot_feed</b> data packet.
<b>value</b>	A pointer to a float variable. A successful call will return the decoded data value of the peer statistic. Before the value is returned, the <b>RSiGetHotItem</b> function: <ul style="list-style-type: none"><li>• Determines the format of the data field as being either <b>SiFloat</b> or <b>SiLong</b> and extracts the data value for further processing.</li><li>• Determines the data value as being either type <b>SiQuantity</b> or type <b>SiCounter</b> and performs one of the actions listed here:<ul style="list-style-type: none"><li>– If the data value is of type <b>SiQuantity</b>, the subroutine returns the <b>val</b> field of the <b>SpmiHotItems</b> structure.</li><li>– If the data value is of type <b>SiCounter</b>, the subroutine returns the value of the <b>val_change</b> field of the <b>SpmiHotItems</b> structure divided by the elapsed number of seconds since the previous time a data value was requested for this set of statistics.</li></ul></li></ul>
<b>absvalue</b>	A pointer to a float variable. A successful call will return the decoded value of the <b>val</b> field of the <b>SpmiHotItems</b> structure of the peer statistic. In case of a statistic of type <b>SiQuantity</b> , this value will be the same as the one returned in the argument <i>value</i> . In case of a peer statistic of type <b>SiCounter</b> , the value returned is the absolute value of the counter.
<b>name</b>	A pointer to a character pointer. A successful call will return a pointer to the name of the peer context for which the data value was read.

## Return Values

The **RSiGetHotItem** subroutine returns a pointer to the current **SpmiHotVals** structure within the hotset. If no more **SpmiHotItems** elements are available, the subroutine returns a NULL value. The structure returned contains the data, such as threshold, which may be relevant for presentation of the results of an **SpmiGetHotSet** subroutine call to end-users. In the returned **SpmiHotVals** structure, all fields contain the correct values as declared, except for the following:

<b>stat</b>	Declared as <b>SpmiStatHdl</b> , actually points to a valid <b>SpmiStat</b> structure. By casting the handle to a pointer to <b>SpmiStat</b> , data in the structure can be accessed.
<b>grandpa items</b>	Contains the <b>cx_handle</b> for the parent context of the peer contexts. When using the <b>Spmi</b> interface this is an array of <b>SpmiHotItems</b> structures. When using the <b>RSiGetHotItem</b> subroutine, the array is empty and attempts to access it will likely result in segmentation faults or access of not valid data.

**path** Will contain the path to the parent of the peer contexts. Even when the peer contexts are multiple levels below the parent context, the path points to the top context because the peer context identifiers in the **SpmiHotItems** elements will contain the path name from there and on. For example, if the hotvals peer set defines all volume groups, the path specified in the returned **SpmiHotVals** structure would be “FS” and the path name in one **SpmiHotItems** element may be “rootvg/lv01”. When combined with the metric name from the **stat** field, the full path name can be constructed as, for example, “FS/rootvg/lv01/%totfree”.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSi error codes are described in *List of RSi Error Codes*.

## Files

**/usr/include/sys/Rsi.h** Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “RSiOpen Subroutine” on page 75
- “RSiCreateHotSet Subroutine” on page 56.

---

## RSiGetRawValue Subroutine

### Purpose

Returns a pointer to a valid **SpmiStatVals** structure for a given **SpmiStatVals** pointer by extraction from a **data\_feed** packet. This subroutine call should only be issued from a callback function after it has been verified that a **data\_feed** packet was received from the host identified by the first argument.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
struct SpmiStatVals RSiGetRawValue(rhandle, svp, index)
RSiHandle rhandle;
struct SpmiStatVals *svp;
int *index;
```

### Description

The **RSiGetRawValue** subroutine performs the following:

1. Finds an **SpmiStatVals** structure in the received data packet based upon the second argument to the subroutine call. This involves a lookup operation in tables maintained internally by the RSi interface.

2. Updates the **struct SpmiStat** pointer in the **SpmiStatVals** structure to point at a valid **SpmiStat** structure.
3. Returns a pointer to the **SpmiStatVals** structure. The returned pointer points to a static area and is only valid until the next execution of **RSiGetRawValue**.
4. Updates an integer variable with the index into the **ValsSet** array of the **data\_feed** packet, which corresponds to the second argument to the call.

This subroutine is part of the Performance Toolbox for AIX licensed product.

## Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**svp** A handle of type **struct SpmiStatVals**, which was previously returned by a successful **RSiPathAddSetStat** (“**RSiPathAddSetStat Subroutine**” on page 78) subroutine call.

**index** A pointer to an integer variable. When the subroutine call succeeds, the index into the **ValsSet** array of the data feed packet is returned. The index corresponds to the element that matches the **svp** argument to the subroutine.

## Return Values

If successful, the subroutine returns a pointer; otherwise NULL is returned and an error text may be placed in the external character array **RSiErrMsg**.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiErrMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiErrMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSI Error Codes*.

## Files

**/usr/include/sys/Rsi.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “**RSiOpen Subroutine**” on page 75
- “**RSiPathAddSetStat Subroutine**” on page 78.

---

## RSiGetValue Subroutine

### Purpose

Returns a data value for a given **SpmiStatVals** pointer by extraction from the **data\_feed** packet. This subroutine call should only be issued from a callback function after it has been verified that a **data\_feed** packet was received from the host identified by the first argument.

## Library

RSI Library (**libSpmi.a**)

## Syntax

```
#include sys/Rsi.h
float RSiGetValue(rhandle, svp)
RSiHandle rhandle;
struct SpmiStatVals *svp;
```

## Description

The **RSiGetValue** subroutine provides the following:

1. Finds an **SpmiStatVals** structure in the received data packet based upon the second argument to the subroutine call. This involves a lookup operation in tables maintained internally by the RSi interface.
2. Determines the format of the data field as being either **SiFloat** or **SiLong** and extracts the data value for further processing based upon its data format.
3. Determines the value as either of type **SiQuantity** or **SiCounter**. If the former is the case, the data value returned is the **val** field in the **SpmiStatVals** structure. If the latter type is found, the value returned by the subroutine is the **val\_change** field divided by the elapsed number of seconds since the previous data packet's time stamp.

This subroutine is part of the Performance Toolbox for AIX licensed product.

## Parameters

**rhandle** Must be an **RSiHandle**, previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**svp** A handle of type **struct SpmiStatVals**, which was previously returned by a successful **RSiPathAddSetStat** (“**RSiPathAddSetStat Subroutine**” on page 78) subroutine call.

## Return Values

If successful, the subroutine returns a non-negative value; otherwise it returns a negative value less than or equal to -1.0. A NULL error text is placed in the external character array **RSiEMsg** regardless of the subroutine's success or failure.

## Error Codes

All RSi subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSi error codes are described in *List of RSi Error Codes*.

## Files

**/usr/include/sys/Rsi.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSi.

## Related Information

For related information, see:

- “RSiOpen Subroutine” on page 75
- “RSiPathAddSetStat Subroutine” on page 78

---

## RSiInit Subroutine

### Purpose

Allocates or changes the table of RSi handles.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
RSiHandle RSiInit(count)
int count;
```

### Description

Before any other **RSi** call is executed, a data-consumer program must issue the **RSiInit** call. Its purpose is to either:

- Allocate an array of **RSiHandleStruct** structures and return the address of the array to the data-consumer program.
- Increase the size of a previously allocated array of **RSiHandleStruct** structures and initialize the new array with the contents of the previous one.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Parameters

**count** Must specify the number of elements in the array of RSi handles. If the call is used to expand a previously allocated array, this argument must be larger than the current number of array elements. It must always be larger than zero. Specify the size of the array to be at least as large as the number of hosts your data-consumer program can talk to at any point in time.

### Return Values

If successful, the subroutine returns the address of the allocated array. If an error occurs, an error text is placed in the external character array **RSiErrMsg** and the subroutine returns NULL. When used to increase the size of a previously allocated array, the subroutine first allocates the new array, then moves the entire old array to the new area. Application programs should, therefore, refer to elements in the RSi handle array by index rather than by address if they anticipate the need for expanding the array. The array only needs to be expanded if the number of remote hosts a data-consumer program talks to might increase over the life of the program.

An application that calls **RSiInit** repeatedly needs to preserve the previous address of the **RSiHandle** array while the **RSiInit** call is re-executed. After the call has completed successfully, the calling program should free the previous array using the **free** subroutine.

### Error Codes

All RSi subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSi error codes are described in *List of RSi Error Codes*.

## Files

`/usr/include/sys/Rsi.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see the “RSiClose Subroutine” on page 55.

---

## RSiInstantiate Subroutine

### Purpose

Creates (instantiates) all subcontexts of an **SpmiCx** context object.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
int RSiInstantiate(rhandle, context)
RSiHandle rhandle;
cx_handle *context;
```

### Description

The **RSiInstantiate** subroutine performs the following actions:

1. Validates that the context identified by the second argument exists.
2. Instantiates the context so that all subcontexts of that context are created in the context hierarchy. Note that this subroutine call currently only makes sense if the context’s **SilnstFreq** is set to **SiContlnst** or **SiCfglnst** because all other contexts would have been instantiated whenever the **xmservd** daemon was started.

The **RSiInstantiate** subroutine explicitly instantiates the subcontexts of an instantiable context. If the context is not instantiable, do not call the **RSiInstantiate** subroutine.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Parameters

**rhandle** Must point to a structure of type **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**context** Must be a handle of type **cx\_handle**, which was previously returned by a successful **RSiPathGetCx** (“**RSiPathGetCx Subroutine**” on page 79) subroutine call.



**resy\_callb** Must be either NULL or a pointer to a function that processes **i\_am\_back** packets as they are received from the **xmservd** daemons on remote hosts for the duration of the **RSiInvite** subroutine call. When the callback function is invoked, it is passed three arguments as described in the following information.

If this argument is specified as NULL, a callback function internal to the **RSiInvite** subroutine receives any **i\_am\_back** packets and uses them to build the table of host names the function returns.

**excp\_callb** Must be NULL or a pointer to a function that processes **except\_rec** packets as they are received from the **xmservd** daemons on remote hosts. If a NULL pointer is passed, your application does not receive **except\_rec** messages. When this callback function is invoked, it is passed three arguments as described in the following information.

This argument always overrides the corresponding argument of any previous **RSiInvite** or **RSiOpen** call, and it can be overridden by subsequent executions of either. In this way, your application can turn exception monitoring on and off. For an **RSiOpen** to override the exception processing specified by a previous open call, the connection must first be closed with the **RSiClose** call. That's because an **RSiOpen** against an already active handle is treated as a no-operation.

The **resy\_callb** and **excp\_callb** functions in your application are called with the following three arguments:

- An **RSiHandle**. The RSi handle pointed to is almost certain not to represent the host that sent the packet. Ignore this argument, and use only the second one: the pointer to the input buffer.
- A pointer of type **pack \*** to the input buffer containing the received packet. Always use this pointer rather than the pointer in the **RSiHandle** structure.
- A pointer of type **struct sockaddr\_in \*** to the IP address of the originating host.

## Return Values

If successful, the subroutine returns an array of character pointers, each of which contains a host name of a host that responded to the invitation. The returned host names are actually constructed as two “words” with the first one being the host name returned by the host in response to an **are\_you\_there** request; the second one being the character form of the host's IP address. The two “words” are separated by one or more blanks. This format is suitable as an argument to the **RSiOpen (“RSiOpen Subroutine” on page 75)** subroutine call. In addition, the external integer variable **RSiInvTabActive** contains the number of host names found. The returned pointer to an array of host names must not be freed by the subroutine call. The calling program should not assume that the pointer returned by this subroutine call remains valid after subsequent calls to **RSiInvite**. If the call is not successful, an error text is placed in the external character array **RSiErrMsg**, an error number is placed in **RSiErrno**, and the subroutine returns NULL.

The list of host names returned by **RSiInvite** does not include the hosts your program has already established a connection with through an **RSiOpen** call. Your program is responsible for keeping track of such hosts. If you need a list of both sets of hosts, either let the **RSiInvite** call be the first one issued from your program or merge the list of host names returned by the call with the list of hosts to which you have connections.

## Error Codes

All RSi subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiErrMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiErrMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSi error codes are described in *List of RSi Error Codes*.

## Files

`/usr/include/sys/Rsi.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see “RSiOpen Subroutine” on page 75.

---

## RSiMainLoop Subroutine

### Purpose

Allows an application to suspend execution and wait to get awakened when data feeds arrive.

### Library

RSI Library (`libSpmi.a`)

### Syntax

```
#include sys/Rsi.h
void RSiMainLoop(msecs)
int msecs;
```

### Description

The **RSiMainLoop** subroutine:

1. Allows the data-consumer program to suspend processing while waiting for **data\_feed** packets to arrive from one or more **xmservd** daemons.
2. Tells the subroutine that waits for data feeds to return control to the data-consumer program so that the latter can check for and react to other events.
3. Invokes the subroutine to process **data\_feed** packets for each such packet received.

To work properly, the **RSiMainLoop** subroutine requires that at least one **RSiOpen** (“**RSiOpen Subroutine**” on page 75) call has been successfully completed and that the connection has not been closed.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Parameters

**msecs** The minimum elapsed time in milliseconds that the subroutine should continue to attempt receives before returning to the caller. Notice that your program releases control for as many milliseconds you specify but that the callback functions defined on the **RSiOpen** call may be called repetitively during that time.

### Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSi Error Codes*.

## Files

`/usr/include/sys/Rsi.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see “RSiOpen Subroutine” on page 75.

---

## RSiNextCx Subroutine

### Purpose

Returns the next subcontext of an **SpmiCx** context.

### Library

RSI Library (`libSpmi.a`)

### Syntax

```
#include sys/Rsi.h
struct SpmiCxLink *RSiNextCx(rhandle, context, link, name,
descr)
RSiHandle rhandle;
cx_handle *context;
struct SpmiCxLink *link;
char **name;
char **descr;
```

### Description

The **RSiNextCx** subroutine:

1. Validates that the context identified by the second argument exists.
2. Returns a handle to the next element of the list of subcontexts defined for the context.
3. Returns the short name and description of the subcontext.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Parameters

**rhandle** Must point to a structure of type **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**context** Must be a handle of type **cx\_handle**, which was previously returned by a successful **RSiPathGetCx** (“**RSiPathGetCx Subroutine**” on page 79) subroutine call.

**link** Must be a pointer to a structure of type **struct SpmiCxLink**, which was previously returned by a successful **RSiFirstCx** (“**RSiFirstCx Subroutine**” on page 60) or **RSiNextCx** subroutine call.

**name** Must be a pointer to a pointer to a character array. The pointer must be initialized to point at a character array pointer. When the subroutine call is successful, the short name of the subcontext is returned in the character array pointer.

**descr** Must be a pointer to a pointer to a character array. The pointer must be initialized to point at a character array pointer. When the subroutine call is successful, the description of the subcontext is returned in the character array pointer.

## Return Values

If successful, the subroutine returns a pointer to a structure of type **struct SpmiCxLink**. If an error occurs, or if no more subcontexts exist for the context, NULL is returned and an error text may be placed in the external character array **RSiEMsg**.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSi error codes are described in *List of RSi Error Codes*.

## Files

`/usr/include/sys/Rsi.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “RSiFirstCx Subroutine” on page 60
- “RSiOpen Subroutine” on page 75
- “RSiPathGetCx Subroutine” on page 79.

---

## RSiNextStat Subroutine

### Purpose

Returns the next statistic of an **SpmiCx** context.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h

struct SpmiStatLink *RSiNextStat(rhandle, context, link, name,
descr)
RSiHandle rhandle;
cx_handle *context;
struct SpmiStatLink *link;
char **name;
char **descr;
```

### Description

The **RSiNextStat** subroutine:

1. Validates that a context identified by the second argument exists.
2. Returns a handle to the next element of the list of statistics defined for the context.
3. Returns the short name and description of the statistic.

This subroutine is part of the Performance Toolbox for AIX licensed product.

## Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**”) subroutine.

**context** Must be a handle of type **cx\_handle**, which was previously returned by a successful **RSiPathGetCx** (“**RSiPathGetCx Subroutine**” on page 79) subroutine call.

**link** Must be a pointer to a structure of type **struct SpmiStatLink**, which was previously returned by a successful **RSiFirstStat** (“**RSiFirstStat Subroutine**” on page 62) or **RSiNextStat** subroutine call.

**name** Must be a pointer to a pointer to a character array. The pointer must be initialized to point at a character array pointer. When the subroutine call is successful, the short name of the statistics value is returned in the character array pointer.

**descr** Must be a pointer to a pointer to a character array. The pointer must be initialized to point at a character array pointer. When the subroutine call is successful, the description of the statistics value is returned in the character array pointer.

## Return Values

If successful, the subroutine returns a pointer to a structure of type **struct SpmiStatLink**. If an error occurs, or if no more statistics exists for the context, NULL is returned and an error text may be placed in the external character array **RSiErrMsg**.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiErrMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiErrMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSI Error Codes*.

## Files

**/usr/include/sys/Rsi.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “**RSiFirstStat Subroutine**” on page 62
- “**RSiOpen Subroutine**”
- “**RSiPathGetCx Subroutine**” on page 79.

---

## RSiOpen Subroutine

### Purpose

Initializes the RSI interface for a remote host.

## Library

RSI Library (**libSpmi.a**)

## Syntax

```
#include sys/Rsi.h
int RSiOpen(rhandle, wait, bufsize, hostID, feed_callb,
           resy_callb, excp_callb)
RSiHandle rhandle;
int wait;
int bufsize;
char *hostID;
int (*feed_callb)();
int (*resy_callb)();
int (*excp_callb)();
```

## Description

The **RSiOpen** subroutine performs the following actions:

1. Establishes the issuing data-consumer program as a data consumer known to the **xmservd** daemon on a particular host. The subroutine does this by sending an **are\_you\_there** packet to the host.
2. Initializes an RSi handle for subsequent use by the data-consumer program.

This subroutine is part of the Performance Toolbox for AIX licensed product.

## Parameters

The arguments to the subroutine are:

**rhandle** Must point to an element of the **RSiHandleStruct** array, which is returned by a previous **RSiInit** (“**RSiInit Subroutine**” on page 68) call. If the subroutine is successful the structure is initialized and ready to use as a handle for subsequent RSi interface subroutine calls.

**wait** Must specify the timeout in milliseconds that the RSi interface shall wait for a response when using the request-response functions. On LANs, a reasonable value for this argument is 100 milliseconds. If the response is not received after the specified wait time, the library subroutines retry the receive operation until five times the wait time has elapsed before returning a timeout indication. The wait time must be zero or more milliseconds.

**bufsize** Specifies the maximum buffer size to be used for constructing network packets. This size must be at least 4,096 bytes. The buffer size determines the maximum packet length that can be received by your program and sets the limit for the number of data values that can be received in one **data\_feed** packet. There’s no point in setting the buffer size larger than that of the **xmservd** daemon because both must be able to handle the packets. If you need large sets of values, you can use the command line argument **-b** of **xmservd** to increase its buffer size up to 16,384 bytes.

The fixed part of a **data\_feed** packet is 104 bytes and each value takes 32 bytes. A buffer size of 4,096 bytes allows up to 124 values per packet.

**hostID** Must be a character array containing the identification of the remote host whose **xmservd** daemon is the one with which you want to talk. The first characters of the host identification (up to the first white space) is used as the host name. The full host identification is stored in the **RSiHandle** field **longname** and may contain any description that helps the end user identify the host used. The host name may be either in long format (including domain name) or in short format.

**feed\_callb** Must be a pointer to a function that processes **data\_feed** packets as they are received from the **xmservd** daemon. When this callback function is invoked, it is passed three arguments as described in the following information.

**resy\_callb** Must be a pointer to a function that processes **i\_am\_back** packets as they are received from the **xmservd** daemon. When this callback function is invoked it is passed three arguments as described in the following information.

**excp\_callb** Must be NULL or a pointer to a function that processes **except\_rec** packets as they are received from the **xmservd** daemon. If a NULL pointer is passed, your application does not receive **except\_rec** messages. When this callback function is invoked, it is passed three arguments as described in the following information. This argument always overrides the corresponding argument of any previous **RSiInvite** (“**RSiInvite Subroutine**” on page 70) or **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine call and can itself be overridden by subsequent executions of either. In this way, your application can turn exception monitoring on and off. For an **RSiOpen** call to override the exception processing specified by a previous open call, the connection must first be closed with the **RSiClose** (“**RSiClose Subroutine**” on page 55) subroutine call.

The **feed\_callb**, **resy\_callb**, and **excp\_callb** functions are called with the arguments:

**RSiHandle**. When a **data\_feed** packet is received, the structure pointed to is guaranteed to represent the host sending the packet. In all other situations the **RSiHandle** structure may represent any of the hosts to which your application is talking.

Pointer of type **pack \*** to the input buffer containing the received packet. In callback functions, always use this pointer rather than the pointer in the **RSiHandle** structure.

Pointer of type **struct sockaddr\_in \*** to the IP address of the originating host.

## Return Values

If successful, the subroutine returns zero and initializes the array element of type **RSiHandle** pointed to by **rhandle**. If an error occurs, error text is placed in the external character array **RSiEMsg** and the subroutine returns a negative value.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSI Error Codes*.

## Files

**/usr/include/sys/Rsi.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “**RSiClose Subroutine**” on page 55
- “**RSiInvite Subroutine**” on page 70

---

## RSiPathAddSetStat Subroutine

### Purpose

Add a single statistics value to an already defined **SpmiStatSet**.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
struct SpmiStatVals *RSiPathAddSetStat(rhandle, statset,
path)
RSiHandle rhandle;
struct SpmiStatSet *statset;
char *path;
```

### Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**statset** Must be a pointer to a structure of type **struct SpmiStatSet**, which was previously returned by a successful **RSiCreateStatSet** (“**RSiCreateStatSet Subroutine**” on page 57) subroutine call.

**path** Must be the full value path name of the statistics value to add to the **SpmiStatSet**. The value path name must not include a terminating slash. Note that value path names never start with a slash.

### Return Values

If successful, the subroutine returns a pointer to a structure of type **struct SpmiStatVals**. If an error occurs, NULL is returned and an error text may be placed in the external character array **RSiEMsg**. If you attempt to add more values to a statset than the current local buffer size allows, **RSiErrno** is set to **RSiTooMany**. If you attempt to add more values than the buffer size of the remote host’s **xmservd** daemon allows, **RSiErrno** is set to **RSiBadStat** and the status field in the returned packet is set to **too\_many\_values**.

The external integer **RSiMaxValues** holds the maximum number of values acceptable with the data-consumer’s buffer size.

### Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSi error codes are described in *List of RSi Error Codes*.

### Files

**/usr/include/sys/Rsi.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “RSiCreateStatSet Subroutine” on page 57
- “RSiOpen Subroutine” on page 75.

---

## RSiPathGetCx Subroutine

### Purpose

Searches the context hierarchy for an **SpmiCx** context that matches a context path name.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
cx_handle *RSiPathGetCx(rhandle, path)
RSiHandle rhandle;
char *path;
```

### Description

The **RSiPathGetCx** subroutine performs the following actions:

1. Searches the context hierarchy for a given path name of a context.
2. Returns a handle to be used when subsequently referencing the context.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**path** A path name of a context for which a handle is to be returned. The context path name must be the full path name and must not include a terminating slash. Note that context path names never start with a slash.

### Return Values

If successful, the subroutine returns a handle defined as a pointer to a structure of type **cx\_handle**. If an error occurs, NULL is returned and an error text may be placed in the external character array **RSiEMsg**.

### Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSI Error Codes*.

## Files

`/usr/include/sys/Rsi.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “RSiFirstCx Subroutine” on page 60
- “RSiOpen Subroutine” on page 75
- “RSiNextCx Subroutine” on page 73.

---

## RSiStartFeed Subroutine

### Purpose

Tells **xmservd** to start sending data feeds for a **statset**.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
int RSiStartFeed(rhandle, statset, msecs)
RSiHandle rhandle;
struct SpmiStatSet *statset;
int msecs;
```

### Description

The **RSiStartFeed** subroutine performs the following function:

1. Informs **xmservd** of the frequency with which it is required to send **data\_feed** packets.
2. Tells the **xmservd** to start sending **data\_feed** packets.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**statset** Must be a pointer to a structure of type **struct SpmiStatSet**, which was previously returned by a successful **RSiCreateStatSet** (“**RSiCreateStatSet Subroutine**” on page 57) subroutine call.

**msecs** The number of milliseconds between the sending of **data\_feed** packets. This number is rounded to a multiple of **min\_remote\_int** milliseconds by the **xmservd** daemon on the remote host. This minimum interval can be modified through the **-i** command line interval to **xmservd**.

### Return Values

If successful, the subroutine returns zero; otherwise it returns -1 and an error text may be placed in the external character array **RSIEMsg**.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- `extern char RSiEMsg[];`
- `extern int RSiErrno;`

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSi Error Codes*.

## Files

`/usr/include/sys/Rsi.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “RSiCreateStatSet Subroutine” on page 57
- “RSiOpen Subroutine” on page 75
- “RSiStopFeed Subroutine” on page 84.

---

## RSiStartHotFeed Subroutine

### Purpose

Tells **xmservd** to start sending hot feeds for a hotset or to start checking for if exceptions or SNMP traps should be generated.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
int RSiStartFeed(rhandle, hotset, msecs)
RSiHandle rhandle;
struct SpmiHotSet *hotset;
int msecs;
```

### Description

The **RSiStartHotFeed** subroutine performs the following function:

1. Informs **xmservd** of the frequency with which it is required to send **hot\_feed** packets, if the hotset is defined to generate **hot\_feed** packets.
2. Informs **xmservd** of the frequency with which it is required to check if exceptions or SNMP traps should be generated. This is only done if it is specified for the hotset that exceptions and/or SNMP traps should be generated.
3. Tells the **xmservd** to start sending **data\_feed** packets and/or start checking for exceptions or traps.

This subroutine is part of the Performance Toolbox for AIX licensed product.

## Parameters

**rhandle** Must be an **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**hotset** Must be a pointer to a structure of type **struc SpmiHotSet**, which was previously returned by a successful **RSiCreateHot** (“**RSiCreateHotSet Subroutine**” on page 56) subroutine call.

**msecs** The number of milliseconds between the sending of **hot\_feed** packets and/or the number of milliseconds between checks for if exceptions or SNMP traps should be generated. This number is rounded to a multiple of **min\_remote\_int** milliseconds by the **xmservd** daemon on the remote host. This minimum interval can be modified through the **-i** command line interval to **xmservd**.

## Return Values

If successful, the subroutine returns zero; otherwise it returns -1 and an error text may be placed in the external character array **RSiEMsg**.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSi Error Codes*.

## Files

**/usr/include/sys/Rsi.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “**RSiCreateHotSet Subroutine**” on page 56
- “**RSiOpen Subroutine**” on page 75
- “**RSiChangeHotFeed Subroutine**” on page 54
- “**RSiStopHotFeed Subroutine**” on page 85.

---

## RSiStatGetPath Subroutine

### Purpose

Finds the full path name of a statistic identified by a **SpmiStatVals** pointer.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
```

```
char *RSiStatGetPath(rhandle, svp)
RSiHandle rhandle;
struct SpmiStatVals *svp;
```

## Description

The **RSiStatGetPath** subroutine performs the following:

1. Validates that the **SpmiStatVals** statistic identified by the second argument does exist.
2. Returns a pointer to a character array containing the full value path name of the statistic.

The memory area pointed to by the returned pointer is freed when the **RSiStatGetPath** subroutine call is repeated. For each invocation of the subroutine, a new memory area is allocated and its address returned.

If the calling program needs the returned character string after issuing the **RSiStatGetPath** subroutine call, the program must copy the returned string to locally allocated memory before reissuing the subroutine call.

This subroutine is part of the Performance Toolbox for AIX licensed product.

## Parameters

**rhandle** Must be an **RSiHandle**, previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**svp** Must be a handle of type **struct SpmiStatVals** as returned by a successful **RSiPathAddSetStat** (“**RSiPathAddSetStat Subroutine**” on page 78) subroutine call.

## Return Values

If successful, the **RSiStatGetPath** subroutine returns a pointer to a character array containing the full path name of the statistic. If unsuccessful, the subroutine returns a NULL value and an error text may be placed in the external character array **RSiErrMsg**.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiErrMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiErrMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSi Error Codes*.

## Files

**/usr/include/sys/Rsi.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “**RSiOpen Subroutine**” on page 75
- “**RSiPathAddSetStat Subroutine**” on page 78.

---

## RSiStopFeed Subroutine

### Purpose

Tells **xmservd** to stop sending data feeds for a **statset**.

### Library

RSI Library (**libSpmi.a**)

### Syntax

```
#include sys/Rsi.h
int RSiStopFeed(rhandle, statset, erase)
RSiHandle rhandle;
struct SpmiStatSet *statset;
boolean erase;
```

### Description

The **RSiStopFeed** subroutine instructs the **xmservd** of a remote system to:

1. Stop sending **data\_feed** packets for a given **SpmiStatSet**. If the daemon is not told to erase the **SpmiStatSet**, feeding of data can be resumed by issuing the **RSiStartFeed** (“**RSiStartFeed Subroutine**” on page 80) subroutine call for the **SpmiStatSet**.
2. Optionally tells the daemon and the API library subroutines to erase all their information about the **SpmiStatSet**. Subsequent references to the erased **SpmiStatSet** are not valid.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Parameters

**rhandle** Must point to a structure of type **RSiHandle**, which was previously initialized by the **RSiOpen** (“**RSiOpen Subroutine**” on page 75) subroutine.

**statset** Must be a pointer to a structure of type **struct SpmiStatSet**, which was previously returned by a successful **RSiCreateStatSet** (“**RSiCreateStatSet Subroutine**” on page 57) subroutine call. Data feeding must have been started for this **SpmiStatSet** via a previous **RSiStartFeed** (“**RSiStartFeed Subroutine**” on page 80) subroutine call.

**erase** If this argument is set to true, the **xmservd** daemon on the remote host discards all information about the named **SpmiStatSet**. Otherwise the daemon maintains its definition of the set of statistics.

### Return Values

If successful, the subroutine returns zero, otherwise -1. A NULL error text is placed in the external character array **RSiEMsg** regardless of the subroutine’s success or failure.

### Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSI error codes are described in *List of RSI Error Codes*.

## Files

`/usr/include/sys/Rsi.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “RSiOpen Subroutine” on page 75
- “RSiStartFeed Subroutine” on page 80.

---

## RSiStopHotFeed Subroutine

### Purpose

Tells `xmservd` to stop sending hot feeds for a hotset and to stop checking for exception and SNMP trap generation.

### Library

RSI Library (`libSpmi.a`)

### Syntax

```
#include sys/Rsi.h
int RSiStopFeed(rhandle, hotset, erase)
RSiHandle rhandle;
struct SpmiHotSet *hotset;
boolean erase;
```

### Description

The `RSiStopHotFeed` subroutine instructs the `xmservd` of a remote system to:

1. Stop sending **hot\_feed** packets or check if exceptions or SNMP traps should be generated for a given **SpmiHotSet**. If the daemon is not told to erase the **SpmiHotSet**, feeding of data can be resumed by issuing the `RSiStartHotFeed` (“[RSiStartHotFeed Subroutine](#)” on page 81) subroutine call for the **SpmiHotSet**.
2. Optionally tells the daemon and the API library subroutines to erase all their information about the **SpmiHotSet**. Subsequent references to the erased **SpmiHotSet** are not valid.

This subroutine is part of the Performance Toolbox for AIX licensed product.

### Parameters

**rhandle** Must point to a structure of type **RSiHandle**, which was previously initialized by the `RSiOpen` (“[RSiOpen Subroutine](#)” on page 75) subroutine.

**hotset** Must be a pointer to a structure of type **struct SpmiHotSet**, which was previously returned by a successful `RSiCreateHotSet` (“[RSiCreateHotSet Subroutine](#)” on page 56) subroutine call. Data feeding must have been started for this **SpmiStatSet** via a previous `RSiStartHotFeed` (“[RSiStartHotFeed Subroutine](#)” on page 81) subroutine call.

**erase** If this argument is set to true, the `xmservd` daemon on the remote host discards all information about the named **SpmiHotSet**. Otherwise the daemon maintains its definition of the set of statistics.

## Return Values

If successful, the subroutine returns zero, otherwise -1. A NULL error text is placed in the external character array **RSiEMsg** regardless of the subroutine's success or failure.

## Error Codes

All RSI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char RSiEMsg[];
- extern int RSiErrno;

If the subroutine returns without an error, the **RSiErrno** variable is set to **RSiOkay** and the **RSiEMsg** character array is empty. If an error is detected, the **RSiErrno** variable returns an error code, as defined in the enum **RSiErrorType**. RSi error codes are described in *List of RSi Error Codes*.

## Files

**/usr/include/sys/Rsi.h** Declares the subroutines, data structures, handles, and macros that an application program can use to access the RSI.

## Related Information

For related information, see:

- “RSiOpen Subroutine” on page 75
- “RSiStartHotFeed Subroutine” on page 81
- “RSiChangeHotFeed Subroutine” on page 54.

---

## rs\_alloc Subroutine

### Purpose

Allocates a resource set and returns its handle.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
rsethandle_t rs_alloc (flags)
unsigned int flags;
```

### Description

The **rs\_alloc** subroutine allocates a resource set and initializes it according to the information specified by the *flags* parameter. The value of the *flags* parameter determines how the new resource set is initialized.

The handle for the new resource set is returned by the subroutine.

## Parameters

- flags* Specifies how the new resource set is initialized. It takes one of the following values, defined in **rset.h**:
- **RS\_EMPTY** (or 0 value): The resource set is initialized to contain no resources.
  - **RS\_SYSTEM**: The resource set is initialized to contain available system resources.
  - **RS\_ALL**: The resource set is initialized to contain all resources.
  - **RS\_PARTITION**: The resource set is initialized to contain the resources in the caller's process partition resource set.

## Return Values

On successful completion, a resource set handle for the new resource set is returned. Otherwise, a value of 0 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **rs\_alloc** subroutine is unsuccessful if one or more of the following are true:

- EINVAL** The *flags* parameter contains an invalid value.  
**ENOMEM** There is not enough space to create the data structures related to the resource set.

## Related Information

“rs\_free Subroutine” on page 88, “rs\_getinfo Subroutine” on page 90, and “rs\_init Subroutine” on page 96.

---

## rs\_discardname Subroutine

### Purpose

Discards a resource set definition from the system resource set registry.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
int rs_discardname(namespace, rname)
char *namespace, *rname;
```

### Description

The **rs\_discardname** subroutine discards from the system global repository the definition of the resource set. The resource set is identified by the *namespace* and *rname* parameters. The specified resource set is removed from the registry, and can no longer be shared with other applications.

In order to be able to discard a name from the global repository, the calling process must have root authority or attachment privilege, and an effective user ID equal to that of the *rname* parameter's creator.

## Parameters

- namespace* Points to a null terminated string corresponding to the name space within which *rname* should be found.
- rname* Points to a null terminated string corresponding to the name of a registered resource set to be discarded.

## Return Values

If successful, a value of 0 is returned. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **rs\_discardname** subroutine is unsuccessful if one or more of the following are true:

- EINVAL** One of the following is true:
- The *rsname* parameter contains a null value.
  - The *namespace* parameter contains a null value.
  - The *rsname* or *namespace* parameters point to an invalid name.
  - The name length is null or greater than the **RSET\_NAME\_SIZE** constant (defined in **rset.h**), or the name contains invalid characters.
- EPERM** One of the following is true:
- The calling process has neither root authority nor **CAP\_NUMA\_ATTACH** privileges.
  - The calling process has neither the same user ID as the creator of the *rsname* definition nor root authority .
  - The *namespace* parameter starts with *sys*. This name space is reserved for system use.
- EFAULT** Invalid address, and/or exceptions outside **errno** range.

## Related Information

“rs\_getnameattr Subroutine” on page 91, “rs\_registername Subroutine” on page 99, and “rs\_getnamedrset Subroutine” on page 92.

---

## rs\_free Subroutine

### Purpose

Frees a resource set.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
void rs_free(rset)
rsethandle_t rset;
```

### Description

The **rs\_free** subroutine frees a resource set identified by the *rset* parameter. The resource set must have been allocated by the **rs\_alloc** subroutine

### Parameters

*rset* Specifies the resource set whose memory will be freed.

## Related Information

The “rs\_alloc Subroutine” on page 86.

---

## rs\_getassociativity Subroutine

### Purpose

Gets the hardware associativity values for a resource.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
int rs_getassociativity (type, id, assoc_array, array_size)
unsigned int type;
unsigned int id;
unsigned int *assoc_array;
unsigned int array_size;
```

### Description

The **rs\_getassociativity** subroutine returns the array of hardware associativity values for a specified resource.

This is a special purpose subroutine intended for specialized root applications needing the hardware associativity value information. The **rs\_getinfo**, **rs\_getrad**, and **rs\_numrads** subroutines are provided for non-root applications to discover system hardware topology.

The calling process must have root authority to get hardware associativity values.

### Parameters

<i>type</i>	Specifies the resource type whose associativity values are requested. The only value supported to retrieve values for a processor is R_PROCS.
<i>id</i>	Specifies the logical resource id whose associativity values are requested.
<i>assoc_array</i>	Specifies the address of an array of unsigned integers to receive the associativity values.
<i>array_size</i>	Specifies the number of unsigned integers in <i>assoc_array</i> .

### Return Values

If successful, a value of 0 is returned. The *assoc\_array* parameter array contains the resource's associativity values. The first entry in the array indicates the number of associativity values returned. If the hardware system does not provide system topology data, a value of 0 is returned in the first array entry. If unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **rs\_getassociativity** subroutine is unsuccessful if one or more of the following are true:

<b>EINVAL</b>	One of the following occurred: <ul style="list-style-type: none"><li>• The <i>array_size</i> parameter was specified as 0.</li><li>• An invalid <i>type</i> parameter was specified.</li></ul>
<b>ENODEV</b>	The resource specified by the <i>id</i> parameter does not exist.
<b>EFAULT</b>	Invalid address.
<b>EPERM</b>	The calling process does not have root authority.

## Related Information

“rs\_getinfo Subroutine”, “rs\_getrad Subroutine” on page 94, and “rs\_numrads Subroutine” on page 96.

---

## rs\_getinfo Subroutine

### Purpose

Gets information about a resource set.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
int rs_getinfo(rset, info_type, flags)
rsethandle_t rset;
rsinfo_t info_type;
unsigned int flags;
```

### Description

The **rs\_getinfo** subroutine retrieves information about the resource set identified by the *rset* parameter. Depending on the value of the *info\_type* parameter, the **rs\_getinfo** subroutine returns information about the number of available processors, the number of available memory pools, or the amount of available memory contained in the resource *rset*. The subroutine can also return global system information such as the maximum system detail level, the symmetric multiprocessor (SMP) and multiple chip module (MCM) system detail levels, and the maximum number of processor or memory pool resources in a resource set.

### Parameters

<i>rset</i>	Specifies a resource set handle of a resource set the information should be retrieved from. This parameter is not meaningful if the <i>info_type</i> parameter is R_MAXSDL, R_MAXPROCS, R_MAXMEMPS, R_SMPSDL, or R_MCMSDL.
<i>info_type</i>	Specifies the type of information being requested. One of the following values (defined in <b>rset.h</b> ) can be used: <ul style="list-style-type: none"><li>• <b>R_NUMPROCS</b>: The number of available processors in the resource set is returned.</li><li>• <b>R_NUMMEMPS</b>: The number of available memory pools in the resource set is returned.</li><li>• <b>R_MEMSIZE</b>: The amount of available memory (in MB) contained in the resource set is returned.</li><li>• <b>R_MAXSDL</b>: The maximum system detail level of the system is returned.</li><li>• <b>R_MAXPROCS</b>: The maximum number of processors that may be contained in a resource set is returned.</li><li>• <b>R_MAXMEMPS</b>: The maximum number of memory pools that may be contained in a resource set is returned.</li><li>• <b>R_SMPSDL</b>: The system detail level that corresponds to the traditional notion of an SMP is returned. A system detail level of 0 is returned if the hardware system does not provide system topology data.</li><li>• <b>R_MCMSDL</b>: The system detail level that corresponds to resources packaged in an MCM is returned. A system detail level of 0 is returned if the hardware system does not have MCMs or does not provide system topology data.</li></ul>
<i>flags</i>	Reserved for future use. Specify as 0.

## Return Values

If successful, the requested information is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **rs\_getinfo** subroutine is unsuccessful if one or more of the following are true:

<b>EINVAL</b>	One of the following is true: <ul style="list-style-type: none"><li>• The <i>info_type</i> parameter specifies an invalid resource type value.</li><li>• The <i>flags</i> parameter was not specified as 0.</li></ul>
<b>EFAULT</b>	Invalid address.

## Related Information

The “rs\_numrads Subroutine” on page 96.

---

## rs\_getnameattr Subroutine

### Purpose

Retrieves the access control information of a resource set definition in the system resource set registry.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
int rs_getnameattr(namespace, rname, attr)
char *namespace, *rname;
rs_attributes_t *attr;
```

### Description

The **rs\_getnameattr** subroutine retrieves from the system resource set registry the access control information of the resource set definition specified by the *namespace* and *rname* parameters.

The owner ID, group ID, and access control information of the specified resource set are stored in the structure pointed to by the *attr* parameter.

**Note:** No special authority or access permission is required to query this information.

### Parameters

<i>namespace</i>	Points to a null terminated string corresponding to the name space within which the <i>rname</i> parameter should be found.
<i>rname</i>	Points to a null terminated string corresponding to the name the information should be retrieved for.

*attr* Points to an **rs\_attributes\_t** structure containing the *owner*, *group*, and *mode* fields, which will be filled by the subroutine. The *mode* field in the **rs\_attributes\_t** structure is used to store the access permissions, and is constructed by logically ORing one or more of the following values, defined in **rset.h**:

- **RS\_IRUSR**: Gives read rights to the name's owner.
- **RS\_IWUSR**: Gives write rights to the name's owner.
- **RS\_IRGRP**: Gives read rights to users of the same group as the name's owner.
- **RS\_IWGRP**: Gives write rights to users of the same group as the name's owner.
- **RS\_IROTH**: Gives read rights to others.
- **RS\_IWOTH**: Gives write rights to others.

Read privilege for a user means that the user can retrieve a resource set definition by issuing a call to the **rs\_getnamedrset** subroutine. Write privilege for a user means that the user can redefine a name by issuing another call to the **rs\_getnamedrset** subroutine.

## Return Values

If successful, a value of 0 is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **rs\_getnameattr** subroutine is unsuccessful if one or more of the following are true:

<b>EINVAL</b>	If one of the following is true: <ul style="list-style-type: none"><li>• The <i>rsname</i> parameter is a null pointer.</li><li>• The <i>namespace</i> parameter is a null pointer.</li><li>• The <i>rsname</i> or <i>namespace</i> parameters point to an invalid name. The name length is 0 or greater than the <b>RSET_NAME_SIZE</b> constant (defined in <b>rset.h</b>), or the <i>rsname</i> parameter contains invalid characters.</li></ul>
<b>ENOENT</b>	The <i>rsname</i> parameter could not be found in the name space identified by the <i>namespace</i> parameter.
<b>EFAULT</b>	Invalid address.

## Related Information

“rs\_registername Subroutine” on page 99, “rs\_discardname Subroutine” on page 87, and “rs\_getnamedrset Subroutine”.

---

## rs\_getnamedrset Subroutine

### Purpose

Retrieves the contents of a named resource set from the system resource set registry.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
int rs_getnamedrset (namespace, rsname, rset)
char *namespace, *rsname;
```

## Description

The **rs\_getnamedrset** subroutine retrieves a resource set definition from the system registry. The *namespace* and *rsname* parameters identify the resource set to be retrieved. The *rset* parameter identifies where the retrieved resource set should be returned. The *namespace* and *rsname* parameters identify a previously registered resource set definition.

The calling process must have root authority or read access rights to the resource set definition in order to retrieve it.

The *rset* parameter must be allocated (using the **rs\_alloc** subroutine) prior to calling the **rs\_getnamedrset** subroutine.

## Parameters

<i>namespace</i>	Points to a null-terminated string corresponding to the name space within which <i>rsname</i> is found.
<i>rsname</i>	Points to a null-terminated string corresponding to the previously registered name of a resource set.
<i>rset</i>	Specifies the resource set handle for the resource set that the registered resource set is copied into. The registered resource set is specified by the <i>rsname</i> parameter.

## Return Values

If successful, a value of 0 is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **rs\_getnamedrset** subroutine is unsuccessful if one or more of the following are true:

<b>EINVAL</b>	One of the following is true: <ul style="list-style-type: none"><li>• The <i>rsname</i> parameter is a null pointer.</li><li>• The <i>namespace</i> parameter is a null pointer.</li><li>• The <i>rsname</i> or <i>namespace</i> parameters point to an invalid name. The name length is 0 or greater than the <b>RSET_NAME_SIZE</b> constant (defined in <b>rset.h</b>), or the <i>rsname</i> parameter contains invalid characters.</li></ul>
<b>ENOENT</b>	The <i>rsname</i> parameter could not be found in the name space identified by the <i>namespace</i> parameter.
<b>EPERM</b>	The calling process has neither read permission on <i>rsname</i> nor root authority.
<b>EFAULT</b>	Invalid address.

## Related Information

“rs\_alloc Subroutine” on page 86, “rs\_registername Subroutine” on page 99, “rs\_getnameattr Subroutine” on page 91, and “rs\_discardname Subroutine” on page 87.

---

## rs\_getpartition Subroutine

### Purpose

Gets the partition resource set to which a process is attached.

### Library

Standard C library (**libc.a**)

## Syntax

```
#include <sys/rset.h>
int rs_getpartition (pid, rset)
pid_t pid;
rsethandle_t rset;
```

## Description

The **rs\_getpartition** subroutine returns the partition resource set attached to the specified process. A process ID value of RS\_MYSELF indicates the partition resource set attached to the current process is requested.

The return value from the **rs\_getpartition** subroutine indicates the type of resource set returned.

A value of RS\_PARTITION\_RSET indicates the process has a partition resource set that is set explicitly. This may be set with the **rs\_setpartition** subroutine or through the use of WLM work classes with resource sets.

A value of RS\_DEFAULT\_RSET indicates the process did not have an explicitly set partition resource set. The system default resource set is returned.

## Parameters

*pid* Specifies the process ID whose partition *rset* is requested.  
*rset* Specifies the resource set to receive the process' partition resource set.

## Return Values

If successful, a value of RS\_PARTITION\_RSET, or RS\_DEFAULT\_RSET is returned. If unsuccessful, a value of -1 is returned and the global **errno** variable is set to indicate the error.

## Error Codes

The **rs\_getpartition** subroutine is unsuccessful if one or more of the following are true:

**EFAULT** Invalid address.  
**ESRCH** The process identified by the *pid* parameter does not exist.

## Related Information

The “ra\_getrset Subroutine” on page 15.

---

## rs\_getrad Subroutine

### Purpose

Returns a system resource allocation domain (RAD) contained in an input resource set.

### Library

Standard C library (**libc.a**)

## Syntax

```
#include <sys/rset.h>
int rs_getrad (rset, rad, sdl, index, flags)
rsethandle_t rset, rad;
unsigned int sdl;
unsigned int index;
unsigned int flags;
```

## Description

The **rs\_getrad** subroutine returns a system RAD at a specified system detail level and index that is contained in an input resource set. If only some of the resources in the specified system RAD are contained in the input resource set, only the resources in both the system RAD and the input resource set are returned.

The input resource set is specified by the *rset* parameter. The output system RAD is identified by the *rad* parameter.

The system RAD is specified by system detail level *sdl* and index number *index*. If only a portion of the specified RAD is contained in *rset*, only that portion is returned in *rad*.

The *rset* and *rad* parameters must be allocated (using the **rs\_alloc** subroutine) prior to calling the **rs\_getrad** subroutine.

## Parameters

<i>rset</i>	Specifies a resource set handle for the input resource set.
<i>rad</i>	Specifies a resource set handle to receive the desired system RAD (contained in the <i>rset</i> parameter).
<i>sdl</i>	Specifies the system detail level of the desired system RAD.
<i>index</i>	Specifies the index of the system RAD that should be returned from among those at the specified <i>sdl</i> . This parameter must belong to the $[0, \text{rs\_numrads}(rset, sdl, flags) - 1]$ interval.
<i>flags</i>	The following flags (defined in <b>rset.h</b> ) can be used to modify the default behavior of the <b>rs_getrad</b> subroutine. By default, the <b>rs_getrad</b> subroutine empties the resource set specified by <i>rad</i> before the specified RAD is retrieved. <ul style="list-style-type: none"><li>• <b>RS_UNION</b>: Instead of emptying <i>rad</i> before the specified RAD is retrieved, the RAD retrieved is added to the contents of <i>rad</i>. On completion, <i>rad</i> contains the union of its original contents and the specified RAD.</li><li>• <b>RS_EXCLUSION</b>: Instead of emptying <i>rad</i> before the specified RAD is retrieved, the resources in the specified RAD that are also in <i>rad</i> are removed from <i>rad</i>. On return, <i>rad</i> contains all the resources it originally contained except those in the specified RAD.</li></ul>

## Return Values

If successful, a value of 0 is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **rs\_getrad** subroutine is unsuccessful if one or more of the following are true:

<b>EINVAL</b>	One of the following is true: <ul style="list-style-type: none"><li>• The <i>flags</i> parameter contains an invalid value.</li><li>• The <i>sdl</i> parameter is greater than the maximum system detail level.</li><li>• The RAD specified by the <i>index</i> parameter does not exist at the system detail level specified by the <i>sdl</i> parameter.</li></ul>
<b>EFAULT</b>	Invalid address.

## Related Information

“rs\_numrads Subroutine”, “rs\_getinfo Subroutine” on page 90, and “rs\_alloc Subroutine” on page 86.

---

## rs\_init Subroutine

### Purpose

Initializes a previously allocated resource set.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
int rs_init (rset, flags)
rsethandle_t rset;
unsigned int flags;
```

### Description

The **rs\_init** subroutine initializes a previously allocated resource set. The resource set is initialized according to information specified by the *flags* parameter.

### Parameters

*rset* Specifies the handle of the resource set to initialize.

*flags* Specifies how the resource set is initialized. It takes one of the following values, defined in **rset.h**:

- **RS\_EMPTY**: The resource set is initialized to contain no resources.
- **RS\_SYSTEM**: The resource set is initialized to contain available system resources.
- **RS\_ALL**: The resource set is initialized to contain all resources.
- **RS\_PARTITION**: The resource set is initialized to contain the resources in the caller’s process partition resource set.

### Return Values

If successful, a value of 0 is returned. If unsuccessful, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

### Error Codes

The **rs\_init** subroutine is unsuccessful if one or more of the following are true:

**EINVAL** The *flags* parameter contains an invalid value.

## Related Information

The “rs\_alloc Subroutine” on page 86.

---

## rs\_numrads Subroutine

### Purpose

Returns the number of system resource allocation domains (RADs) that have available resources.

## Library

Standard C library (**libc.a**)

## Syntax

```
#include <sys/rset.h>
int rs_numrads(rset, sdl, flags)
rsethandle_t rset;
unsigned int sdl;
unsigned int flags;
```

## Description

The **rs\_numrads** subroutine returns the number of system RADs at system detail level *sdl*, that have available resources contained in the resource set identified by the *rset* parameter.

The number of atomic RADs contained in the *rset* parameter is returned if the *sdl* parameter is equal to the maximum system detail level.

## Parameters

*rset* Specifies the resource set handle for the resource set being queried.  
*sdl* Specifies the system detail level in which the caller is interested.  
*flags* Reserved for future use. Specify as 0.

## Return Values

If successful, the number of available RADs at system detail level *sdl*, that have resources contained in the specified resource set is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **rs\_numrads** subroutine is unsuccessful if one or more of the following are true:

**EINVAL** One of the following is true:

- The *flags* parameter contains an invalid value.
- The *sdl* parameter is greater than the maximum system detail level.

**EFAULT** Invalid address.

## Related Information

“rs\_getrad Subroutine” on page 94, and “rs\_getinfo Subroutine” on page 90.

---

## rs\_op Subroutine

### Purpose

Performs a set of operations on one or two resource sets.

### Library

Standard C library (**libc.a**)

## Syntax

```
#include <sys/rset.h>
int rs_op (command, rset1, rset2, flags, id)
unsigned int command;
rsethandle_t rset1, rset2;
unsigned int flags;
unsigned int id;
```

## Description

The `rs_op` subroutine performs the operation specified by the `command` parameter on resource set `rset1` or both resource sets `rset1` and `rset2`.

## Parameters

<i>command</i>	<p>Specifies the operation to apply to the resource sets identified by <i>rset1</i> and <i>rset2</i>. One of the following values, defined in <b>rset.h</b>, can be used:</p> <ul style="list-style-type: none"><li>• <b>RS_UNION</b>: The resources contained in either <i>rset1</i> or <i>rset2</i> are stored in <i>rset2</i>.</li><li>• <b>RS_INTERSECTION</b>: The resources that are contained in both <i>rset1</i> and <i>rset2</i> are stored in <i>rset2</i>.</li><li>• <b>RS_EXCLUSION</b>: The resources in <i>rset1</i> that are also in <i>rset2</i> are removed from <i>rset2</i>. On completion, <i>rset2</i> contains all the resources that were contained in <i>rset2</i> but were not contained in <i>rset1</i>.</li><li>• <b>RS_COPY</b>: All resources in <i>rset1</i> whose type is <i>flags</i> are stored in <i>rset2</i>. If <i>rset1</i> contains no resources of this type, <i>rset2</i> will be empty. The previous content of <i>rset2</i> is lost, while the content of <i>rset1</i> is unchanged.</li><li>• <b>RS_FIRST</b>: The first resource whose type is <i>flags</i> is retrieved from <i>rset1</i> and stored in <i>rset2</i>. If <i>rset1</i> contains no resources of this type, <i>rset2</i> will be empty.</li><li>• <b>RS_NEXT</b>: The resource from <i>rset1</i> whose type is <i>flags</i> and that follows the resource contained in <i>rset2</i> is retrieved and stored in <i>rset2</i>. If no resource of the appropriate type follows the resource specified in <i>rset2</i>, <i>rset2</i> will be empty.</li><li>• <b>RS_NEXT_WRAP</b>: The resource from <i>rset1</i> whose type is <i>flags</i> and that follows the resource contained in <i>rset2</i> is retrieved and stored in <i>rset2</i>. If no resource of the appropriate type follows the resource specified in <i>rset2</i>, <i>rset2</i> will contain the first resource of this type in <i>rset1</i>.</li><li>• <b>RS_ISEMPTY</b>: Test if resource set <i>rset1</i> is empty.</li><li>• <b>RS_ISEQUAL</b>: Test if resource sets <i>rset1</i> and <i>rset2</i> are equal.</li><li>• <b>RS_ISCONTAINED</b>: Test if all resources in resource set <i>rset1</i> are also contained in resource set <i>rset2</i>.</li><li>• <b>RS_TESTRESOURCE</b>: Test if the resource whose type is <i>flags</i> and index is <i>id</i> is contained in resource set <i>rset1</i>.</li><li>• <b>RS_ADDRESOURCE</b>: Add the resource whose type is <i>flags</i> and index is <i>id</i> to resource set <i>rset1</i>.</li><li>• <b>RS_DELRESOURCE</b>: Delete the resource whose type is <i>flags</i> and index is <i>id</i> from resource set <i>rset1</i>.</li></ul>
<i>rset1</i>	Specifies the resource set handle for the first of the resource sets involved in the <i>command</i> operation.
<i>rset2</i>	Specifies the resource set handle for the second of the resource sets involved in the <i>command</i> operation. This resource set is also used, on return, to store the result of the operation, and its previous content is lost. The <i>rset2</i> parameter is ignored on the <b>RS_ISEMPTY</b> , <b>RS_TESTRESOURCE</b> , <b>RS_ADDRESOURCE</b> , and <b>RS_DELRESOURCE</b> commands.
<i>flags</i>	<p>When combined with the <b>RS_COPY</b> command, the <i>flags</i> parameter specifies the type of the resources that will be copied from <i>rset1</i> to <i>rset2</i>. When combined with an <b>RS_FIRST</b> or an <b>RS_NEXT</b> command, the <i>flags</i> parameter specifies the type of the resource that will be retrieved from <i>rset1</i>. This parameter is constructed by logically ORing one or more of the following values, defined in <b>rset.h</b>:</p> <ul style="list-style-type: none"><li>• <b>R_PROCS</b>: processors</li><li>• <b>R_MEMPS</b>: memory pools</li><li>• <b>R_ALL_RESOURCES</b>: processors and memory pools</li></ul> <p>If none of the above are specified for <i>flags</i>, <b>R_ALL_RESOURCES</b> is assumed.</p>

*id* On the RS\_TESTRESOURCE, RS\_ADDRESOURCE, and RS\_DELRESOURCE commands, the *id* parameter specifies the index of the resource to be tested, added, or deleted. This parameter is ignored on the other commands.

## Return Values

If successful, the commands RS\_ISEMPY, RS\_ISEQUAL, RS\_ISCONTAINED, and RS\_TESTRESOURCE return 0 if the tested condition is not met and 1 if the tested condition is met. All other commands return 0 if successful. If unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **rs\_op** subroutine is unsuccessful if one or more of the following are true:

- EINVAL** If one of the following is true:
- *rset1* identifies an invalid resource set.
  - *rset2* identifies an invalid resource set.
  - *command* identifies an invalid operation.
  - *command* is RS\_NEXT or RS\_NEXT\_WRAP\*, and *rset2* does not contain a single resource.
  - *command* is RS\_NEXT or RS\_NEXT\_WRAP\*, and the single resource contained in *rset2* is not also contained in *rset1*.
  - *flags* identifies an invalid resource type.
  - *id* specifies a resource index that is too large.
- EFAULT** Invalid address.

## Related Information

The “rs\_alloc Subroutine” on page 86.

---

## rs\_registername Subroutine

### Purpose

Registers a resource set definition in the system resource set registry.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
int rs_registername(rset, namespace, rname, mode, command)
rsethandle_t rset;
char *namespace, *rname;
unsigned int mode, command;
```

### Description

The **rs\_registername** subroutine registers in the system resource registry (within the name space identified by *namespace*) the definition of the resource set identified by the *rset* handle. The **rs\_registername** subroutine does this by associating with it the name specified by the null terminated string structure pointed to by *rname*.

If *rname* does not exist, the owner and group IDs of *rname* are set to the caller’s owner and group IDs, and the access control information for *rname* is set according to the *mode* parameter.

If *rsname* already exists, its owner and group IDs and its access control information are left unchanged, and the *mode* parameter is ignored. This name can be shared with any applications to identify a dedicated resource set.

Using the *command* parameter, you can ask to overwrite or not to overwrite the *rsname* parameter's registration if it already exists in the global repository within the name space identified by *namespace*. If *rsname* already exists within the specified name space and the *command* parameter is set to **not overwrite**, an error is reported to the calling process.

#### Notes:

1. Registering a resource set definition can only be done by a process that has root authority or CAP\_NUMA\_ATTACH attachment privilege.
2. Overwriting an existing name's registration can be done only by a process that has root authority or write access to this name.

An application registered resource set definition is non-persistent. It does not persist over a system boot.

Both the *namespace* and *rsname* parameters may contain up to 255 characters. They must begin with an ASCII alphanumeric character. Only the period (.), minus (-), and underscore (\_) characters can be mixed with ASCII alphanumeric characters within these strings. Moreover, the names are case-sensitive, which means there is a difference between uppercase and lowercase letters in resource set names and name spaces.

## Parameters

<i>rset</i>	Specifies a resource set handle of a resource set a name should be registered for.
<i>namespace</i>	Points to a null terminated string corresponding to the name space within which <i>rsname</i> will be registered.
<i>rsname</i>	Points to a null terminated string corresponding to the name registered with the setting of the resource set specified by <i>rset</i> .
<i>mode</i>	Specifies the bit pattern that determines the created name access permissions. It is constructed by logically ORing one or more of the following values, defined in <b>rset.h</b> : <ul style="list-style-type: none"><li>• <b>RS_IRUSR</b>: Gives read rights to the name's owner</li><li>• <b>RS_IWUSR</b>: Gives write rights to the name's owner</li><li>• <b>RS_IRGRP</b>: Gives read rights to users of the same group as the name's owner</li><li>• <b>RS_IWGRP</b>: Gives write rights to users of the same group as the name's owner</li><li>• <b>RS_IROTH</b>: Gives read rights to others</li><li>• <b>RS_IWOTH</b>: Gives write rights to others</li></ul>
<i>command</i>	Read privilege for a user means that the user can retrieve a resource set definition (by issuing a call to the <b>rs_getnamedrset</b> subroutine). Write privilege for a user means that the user can redefine a name (by issuing another call to the <b>rs_getnamedrset</b> subroutine). Specifies whether the <i>rsname</i> parameter's registration should be overwritten if it already exists in the global repository. This parameter takes one of the following values, defined in <b>rset.h</b> : <ul style="list-style-type: none"><li>• <b>RS_REDEFINE</b>: The <i>rsname</i> parameter should be redefined if it already exists in the name space identified by <i>namespace</i>. In such a case, the calling process must have write access to <i>rsname</i>.</li><li>• <b>RS_DEFINE</b>: The <i>rsname</i> parameter should not be redefined if it already exists in the name space identified by <i>namespace</i>. If this happens, an error is reported to the calling process</li></ul>

## Return Values

If successful, a value of 0 is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The `rs_registername` subroutine is unsuccessful if one or more of the following are true:

<b>EINVAL</b>	If one of the following is true: <ul style="list-style-type: none"><li>• <i>rsname</i> is a null pointer.</li><li>• <i>namespace</i> is a null pointer.</li><li>• <i>rsname</i> or <i>namespace</i> points to an invalid name. The name length is 0 or greater than the <code>RSET_NAME_SIZE</code> constant (defined in <code>rset.h</code>), or the name contains invalid characters.</li><li>• <i>mode</i> identifies an invalid access rights value.</li><li>• <i>command</i> identifies an invalid command value.</li></ul>
<b>EEXIST</b>	The <i>command</i> parameter is set to <code>RS_DEFINE</code> and <i>rsname</i> already exists in the global repository within the name space identified by <i>namespace</i> .
<b>ENOMEM</b>	There is not enough space to create the data structures related to the registry of this resource set.
<b>EPERM</b>	If one of the following is true: <ul style="list-style-type: none"><li>• The <i>command</i> parameter is set to <code>RS_REDEFINE</code> and the calling process has neither write access to <i>rsname</i> nor root authority .</li><li>• The calling process has neither the attachment privilege nor root authority.</li><li>• The <i>namespace</i> parameter starts with <code>sys</code>. This name space is reserved for system use.</li></ul>
<b>EFAULT</b>	Invalid address, and/or exceptions outside <b>errno</b> range.

## Related Information

“`rs_getnameattr` Subroutine” on page 91, “`rs_discardname` Subroutine” on page 87, and “`rs_getnamedrset` Subroutine” on page 92.

---

## rs\_setnameattr Subroutine

### Purpose

Sets the access control information of a resource set definition in the system resource set registry.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
int rs_setnameattr (namespace, rsname, command, attr)
char *namespace, *rsname;
unsigned int command;
rs_attributes *attr;
```

### Description

The `rs_setnameattr` subroutine sets (depending on the *command* value) one or more of the owner, group, or access control information of the system registry resource set definition specified by the *namespace* and *rsname* parameters.

The owner ID and/or group ID and/or access control information of the *rsname* parameter must be supplied in the structure pointed to by the *attr* parameter.

### Notes:

1. In order to be able to set the attributes of a name, the calling process must have root authority or the attachment privilege and an effective user ID equal to that of the *rsname* parameter’s owner.

2. Root authority is required to change the resource set definition owner ID, or to set its group ID outside of the caller's list of groups.

## Parameters

<i>namespace</i>	Points to a null terminated string corresponding to the name space within which <i>rsname</i> should be found.
<i>rsname</i>	Points to a null terminated string corresponding to the name the information should be retrieved for.
<i>command</i>	Specifies which attributes should be changed. This parameter is constructed by logically ORing one or more of the following values, defined in <b>rset.h</b> : <ul style="list-style-type: none"><li>• <b>RS_OWNER</b>: Set owner as specified in the <i>owner</i> field of <i>attr</i>.</li><li>• <b>RS_GROUP</b>: Set group as specified in the <i>group</i> field of <i>attr</i>.</li><li>• <b>RS_PERM</b>: Set access control information as specified in the <i>mode</i> field of <i>attr</i>.</li></ul>
<i>attr</i>	Points to an <b>rs_attributes_t</b> structure containing the <i>owner</i> , <i>group</i> and <i>mode</i> fields, which will possibly be used by the subroutine for setting attributes. The <i>mode</i> field is used to store the access permissions, and is constructed by logically ORing one or more of the following values, defined in <b>rset.h</b> : <ul style="list-style-type: none"><li>• <b>RS_IRUSR</b>: Gives read rights to the name's owner</li><li>• <b>RS_IWUSR</b>: Gives write rights to the name's owner</li><li>• <b>RS_IRGRP</b>: Gives read rights to users of the same group as the name's owner</li><li>• <b>RS_IWGRP</b>: Gives write rights to users of the same group as the name's owner</li><li>• <b>RS_IROTH</b>: Gives read rights to the others</li><li>• <b>RS_IWOTH</b>: Gives write rights to the others</li></ul>

## Return Values

If successful, a value of 0 is returned. If unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **rs\_setnameattr** subroutine is unsuccessful if one or more of the following are true:

<b>EINVAL</b>	One of the following is true: <ul style="list-style-type: none"><li>• <i>rsname</i> is a null pointer.</li><li>• <i>namespace</i> is a null pointer.</li><li>• <i>rsname</i> or <i>namespace</i> point to an invalid name. Name length is 0 or greater than the <b>RSET_NAME_SIZE</b> constant (defined in <b>rset.h</b>), or name contains invalid characters.</li><li>• <i>command</i> identifies an invalid command value.</li><li>• <i>command</i> includes <b>RS_PERM</b> and the <i>mode</i> field of <i>attr</i> identifies an invalid access rights value.</li><li>• <i>attr</i> is a null pointer.</li></ul>
<b>EPERM</b>	One of the following is true: <ul style="list-style-type: none"><li>• The calling process has neither <b>CAP_NUMA_ATTACH</b> attachment privilege nor root authority.</li><li>• <i>command</i> includes <b>RS_OWNER</b> and the <i>owner</i> field of <i>attr</i> is different from the caller's user ID and the caller does not have root authority.</li><li>• <i>command</i> includes <b>RS_GROUP</b>, the <i>group</i> field of <i>attr</i> is outside of the caller's list of groups, and caller does not have root authority.</li><li>• The <i>namespace</i> parameter starts with <b>sys</b>. This name space is reserved for system use.</li></ul>
<b>ENOENT</b>	<i>rsname</i> could not be found in the name space identified by <i>namespace</i> .
<b>ENOSPC</b>	Out of file-space blocks.
<b>EFAULT</b>	Invalid address; exceptions outside <b>errno</b> range.
<b>ENOSYS</b>	The <b>rs_setnameattr</b> subroutine is not supported by the system.

## Related Information

The “rs\_getnameattr Subroutine” on page 91.

---

## rs\_setpartition Subroutine

### Purpose

Sets the partition resource set of a process.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/rset.h>
int rs_setpartition(pid, rset, flags)
pid_t pid;
rsethandle_t rset;
unsigned int flags;
```

### Description

The **rs\_setpartition** subroutine sets a process' partition resource set. The subroutine can also be used to remove a process' partition resource set.

The partition resource set limits the threads in a process to running only on the processors contained in the partition resource set.

The work component is an existing process identified by the process ID. A process ID value of RS\_MYSELF indicates the attachment applies to the current process.

The following conditions must be met to set a process' partition resource set:

- The calling process must have root authority.
- The resource set must contain processors that are available in the system.
- The new partition resource set must be equal to, or a superset of the target process' effective resource set.
- The target process must not contain any threads that have bindprocessor bindings to a processor.

### Parameters

*pid* Specifies the process ID of the process whose partition resource set is to be set. A value of RS\_MYSELF indicates the current process' partition resource set should be set.

*rset* Specifies the partition resource set to be set. A value of RS\_DEFAULT indicates the process' partition resource set should be removed.

*flags* Reserved for future use. Specify as 0.

### Return Values

If successful, a value of 0 is returned. If unsuccessful, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

### Error Codes

The **rs\_setpartition** subroutine is unsuccessful if one or more of the following are true:

**EINVAL** The *flags* parameter contains an invalid value.

<b>ENODEV</b>	The resource set specified by the <i>rset</i> parameter does not contain any available processors.
<b>ESRCH</b>	The process identified by the <i>pid</i> parameter does not exist.
<b>EFAULT</b>	Invalid address.
<b>ENOMEM</b>	Memory not available.
<b>EPERM</b>	One of the following is true: <ul style="list-style-type: none"> <li>• The calling process does not have root authority.</li> <li>• The process identified by the <i>pid</i> parameter has one or more threads with a bindprocessor processor binding.</li> <li>• The process identified by the <i>pid</i> parameter has an effective resource set and the new partition resource set identified by the <i>rset</i> parameter does not contain all of the effective resource set's resources.</li> </ul>

## Related Information

“rs\_getpartition Subroutine” on page 93 and “ra\_attachrset Subroutine” on page 9.

---

## rsqrt Subroutine

### Purpose

Computes the reciprocal of the square root of a number.

### Libraries

IEEE Math Library (**libm.a**)

System V Math Library (**libmsaa.a**)

### Syntax

```
#include <math.h>
```

```
double rsqrt(double x)
```

### Description

The **rsqrt** command computes the reciprocal of the square root of a number *x*; that is, 1.0 divided by the square root of *x* (1.0/sqrt(*x*)). On some platforms, using the **rsqrt** subroutine is faster than computing 1.0 / sqrt(*x*). The **rsqrt** subroutine uses the same rounding mode used by the calling program.

When using the **libm.a** library, the **rsqrt** subroutine responds to special values of *x* in the following ways:

- If *x* is NaN, then the **rsqrt** subroutine returns NaN. If *x* is a signaling Nan (NaNs), then the **rsqrt** subroutine returns a quiet NaN and sets the **VX** and **VXSNAN** (signaling NaN invalid operation exception) flags in the FPSCR (Floating-Point Status and Control register) to 1.
- If *x* is +/- 0.0, then the **rsqrt** subroutine returns +/- INF and sets the **ZX** (zero divide exception) flag in the FPSCR to 1.
- If *x* is negative, then the **rsqrt** subroutine returns NaN, sets the **errno** global variable to **EDOM**, and sets the **VX** and **VXSQRT** (square root of negative number invalid operation exception) flags in the FPSCR to 1.

When using the **libmsaa.a** library, the **rsqrt** subroutine responds to special values of *x* in the following ways:

- If *x* is +/- 0.0, then the **rsqrt** subroutine returns +/-HUGE\_VAL and sets the **errno** global variable to **EDOM**. The subroutine invokes the **matherr** subroutine, which prints a message indicating a singularity error to standard error output.

- If  $x$  is negative, then the **rsqrt** subroutine returns 0.0 and sets the **errno** global variable to **EDOM**. The subroutine invokes the **matherr** subroutine, which prints a message indicating a domain error to standard error output.

When compiled with **libmsaa.a**, a program can use the **matherr** subroutine to change these error-handling procedures.

## Parameter

$x$  Specifies a double-precision floating-point value.

## Return Values

Upon successful completion, the **rsqrt** subroutine returns the reciprocal of the square root of  $x$ .

**1.0** If  $x$  is 1.0.  
**+0.0** If  $x$  is +INF.

## Error Codes

When using either the **libm.a** or **libmsaa.a** library, the **rsqrt** subroutine may return the following error code:

**EDOM** The value of  $x$  is negative.

## Related Information

The **matherr** subroutine, **sqrt** or **cbrt** (“sqrt, sqrtf, or sqrtl Subroutine” on page 238) subroutine.

---

## rstat Subroutines

### Purpose

Gets performance data from remote kernels.

### Library

(**librpcsvc.a**)

### Syntax

```
#include <rpcsvc/rstat.h>
rstat (host, statp)
char *host;
struct statstime *statp;
```

### Description

The **rstat** subroutine gathers statistics from remote kernels. These statistics are available on items such as paging, swapping and CPU utilization.

### Parameters

*host* Specifies the name of the machine going to be contacted to obtain statistics found in the *statp* parameter.  
*statp* Contains statistics from *host*.

## Return Values

If successful, the **rstat** subroutine fills in the **statstime** for *host* and returns a value of **0**.

## Files

`/usr/include/rpcsvc/rstat.x`

## Related Information

The **rup** command.

The **rstatd** daemon

---

## scalbln, scalblnf, scalblnl, scalbn, scalbnf, scalbnl, or scalb Subroutine

### Purpose

Computes the exponent using FLT\_RADIX=2.

### Syntax

```
#include <math.h>
```

```
double scalbln (x, n)
double x;
long n;
```

```
float scalblnf (x, n)
float x;
long n;
```

```
long double scalblnl (x, n)
long double x;
long n;
```

```
double scalbn (x, n)
double x;
int n;
```

```
float scalbnf (x, n)
float x;
int n;
```

```
long double scalbnl (x, n)
long double x;
int n;
```

```
double scalb(x, y)
double x, y;
```

### Description

The **scalbln**, **scalblnf**, **scalblnl**, **scalbn**, **scalbnf**, and **scalbnl** subroutines compute  $x * FLT\_RADIX^n$  efficiently, not normally by computing  $FLT\_RADIX^n$  explicitly. For AIX, FLT\_RADIX n=2.

The **scalb** subroutine returns the value of the *x* parameter times 2 to the power of the *y* parameter.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be computed.  
*n* Specifies the value to be computed.

## Return Values

Upon successful completion, the **scalbln**, **scalblnf**, **scalblnl**, **scalbn**, **scalbnf**, and **scalbnl** subroutines return  $x * FLT\_RADIX^n$ .

If the result would cause overflow, a range error occurs and the **scalbln**, **scalblnf**, **scalblnl**, **scalbn**, **scalbnf**, and **scalbnl** subroutines return **±HUGE\_VAL**, **±HUGE\_VALF**, and **±HUGE\_VALL** (according to the sign of *x*) as appropriate for the return type of the function.

If the correct value would cause underflow, and is not representable, a range error may occur, and 0.0 is returned.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$  or  $\pm \text{Inf}$ , *x* is returned.

If *n* is 0, *x* is returned.

If the correct value would cause underflow, and is representable, a range error may occur and the correct value is returned.

## Error Codes

If the correct value would overflow, the **scalb** subroutine returns +/-INF (depending on a negative or positive value of the *x* parameter) and sets **errno** to **ERANGE**.

If the correct value would underflow, the **scalb** subroutine returns a value of 0 and sets **errno** to **ERANGE**.

## Related Information

“remainder, remainderf, or remainderl Subroutine” on page 39

**feclearexcept** Subroutine, **fetestexcept** Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

**math.h** in *AIX 5L Version 5.2 Files Reference*.

---

## scandir or alphasort Subroutine

### Purpose

Scans or sorts directory contents.

### Library

Standard C Library (**libc.a**)

## Syntax

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

int scandir(DirectoryName, NameList, Select, Compare)
char * DirectoryName;
struct dirent * (* NameList [ ]);
int (* Select) (struct dirent *);
int (* Compare)(void *, void *);

int alphasort ( Directory1, Directory2)
void *Directory1, *Directory2;
```

## Description

The **scandir** subroutine reads the directory pointed to by the *DirectoryName* parameter, and then uses the **malloc** subroutine to create an array of pointers to directory entries. The **scandir** subroutine returns the number of entries in the array and, through the *NameList* parameter, a pointer to the array.

The *Select* parameter points to a user-supplied subroutine that is called by the **scandir** subroutine to select which entries to include in the array. The selection routine is passed a pointer to a directory entry and should return a nonzero value for a directory entry that is included in the array. If the *Select* parameter is a null value, all directory entries are included.

The *Compare* parameter points to a user-supplied subroutine. This routine is passed to the **qsort** subroutine to sort the completed array. If the *Compare* parameter is a null value, the array is not sorted. The **alphasort** subroutine provides comparison functions for sorting alphabetically.

The memory allocated to the array can be deallocated by freeing each pointer in the array, and the array itself, with the **free** subroutine.

The **alphasort** subroutine treats *Directory1* and *Directory2* as pointers to **dirent** pointers and alphabetically compares them. This subroutine can be passed as the *Compare* parameter to either the **scandir** subroutine or the **qsort** subroutine, or a user-supplied subroutine can be used.

## Parameters

<i>DirectoryName</i>	Points to the directory name.
<i>NameList</i>	Points to the array of pointers to directory entries.
<i>Select</i>	Points to a user-supplied subroutine that is called by the <b>scandir</b> subroutine to select which entries to include in the array.
<i>Compare</i>	Points to a user-supplied subroutine that sorts the completed array.
<i>Directory1, Directory2</i>	Point to <b>dirent</b> structures.

## Return Values

The **scandir** subroutine returns the value -1 if the directory cannot be opened for reading or if the **malloc** subroutine cannot allocate enough memory to hold all the data structures. If successful, the **scandir** subroutine returns the number of entries found.

The **alphasort** subroutine returns the following values:

<b>Less than 0</b>	The <b>dirent</b> structure pointed to by the <i>Directory1</i> parameter is lexically less than the <b>dirent</b> structure pointed to by the <i>Directory2</i> parameter.
<b>0</b>	The <b>dirent</b> structures pointed to by the <i>Directory1</i> parameter and the <i>Directory2</i> parameter are equal.

Greater than 0

The **dirent** structure pointed to by the *Directory1* parameter is lexically greater than the **dirent** structure pointed to by the *Directory2* parameter.

## Related Information

The **malloc**, **free**, **realloc**, **calloc**, **mallopt**, **mallinfo**, or **alloca** subroutine, **opendir**, **readdir**, **telldir**, **seekdir**, **rewinddir**, or **closedir** subroutine, **qsort** (“qsort Subroutine” on page 1) subroutine.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## scanf, fscanf, sscanf, or wsscanf Subroutine

### Purpose

Converts formatted input.

### Library

Standard C Library (**libc.a**)

or (**libc128.a**)

### Syntax

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

```
int scanf ( Format [, Pointer, ... ] )  
const char *Format;
```

```
int fscanf ( Stream, Format [, Pointer, ... ] )  
FILE * Stream;  
const char *Format;
```

```
int sscanf ( String, Format [, Pointer, ... ] )  
const char * String, *Format;
```

```
int wsscanf ( wcs, Format [, Pointer, ... ] )  
const wchar_t * wcs  
const char *Format;
```

### Description

The **scanf**, **fscanf**, **sscanf**, and **wsscanf** subroutines read character data, interpret it according to a format, and store the converted results into specified memory locations. If the subroutine receives insufficient arguments for the format, the results are unreliable. If the format is exhausted while arguments remain, the subroutine evaluates the excess arguments but otherwise ignores them.

These subroutines read their input from the following sources:

<b>scanf</b>	Reads from standard input ( <b>stdin</b> ).
<b>fscanf</b>	Reads from the <i>Stream</i> parameter.
<b>sscanf</b>	Reads from the character string specified by the <i>String</i> parameter.
<b>wsscanf</b>	Reads from the wide character string specified by the <i>wcs</i> parameter.

The **scanf**, **fscanf**, **sscanf**, and **wscanf** subroutines can detect a language-dependent radix character, defined in the program's locale (**LC\_NUMERIC**), in the input string. In the C locale, or in a locale that does not define the radix character, the default radix character is a full stop . (period).

## Parameters

- wcs* Specifies the wide-character string to be read.
- Stream* Specifies the input stream.
- String* Specifies input to be read.
- Pointer* Specifies where to store the interpreted data.

*Format* Contains conversion specifications used to interpret the input. If there are insufficient arguments for the *Format* parameter, the results are unreliable. If the *Format* parameter is exhausted while arguments remain, the excess arguments are evaluated as always but are otherwise ignored.

The *Format* parameter can contain the following:

- Space characters (blank, tab, new-line, vertical-tab, or form-feed characters) that, except in the following two cases, read the input up to the next nonwhite space character. Unless a match in the control string exists, trailing white space (including a new-line character) is not read.
- Any character except a % (percent sign), which must match the next character of the input stream.
- A conversion specification that directs the conversion of the next input field. The conversion specification consists of the following:
  - The % (percent sign) or the character sequence %n\$.

**Note:** The %n\$ character sequence is an X/Open numbered argument specifier. Guidelines for use of the %n% specifier are:

- The value of *n* in %n\$ must be a decimal number without leading 0's and must be in the range from 1 to the **NL\_ARGMAX** value, inclusive. See the **limits.h** file for more information about the **NL\_ARGMAX** value. Using leading 0's (octal numbers) or a larger *n* value can have unpredictable results.
  - Mixing numbered and unnumbered argument specifications in a format string can have unpredictable results. The only exceptions are %% (two percent signs) and %\* (percent sign, asterisk), which can be mixed with the %n\$ form.
  - Referencing numbered arguments in the argument list from the format string more than once can have unpredictable results.
- The optional assignment-suppression character \* (asterisk).
  - An optional decimal integer that specifies the maximum field width.
  - An optional character that sets the size of the receiving variable for some flags. Use the following optional characters:
    - I** Long integer rather than an integer when preceding the **d**, **i**, or **n** conversion codes; unsigned long integer rather than unsigned integer when preceding the **o**, **u**, or **x** conversion codes; double rather than float when preceding the **e**, **f**, or **g** conversion codes.
    - ll** Long long integer rather than an integer when preceding the **d**, **i**, or **n** conversion codes; unsigned long long integer rather than unsigned integer when preceding the **o**, **u**, or **x** conversion codes.
    - L** A long double rather than a float, when preceding the **e**, **f**, or **g** conversion codes; long integer rather than an integer when preceding the **d**, **i**, or **n** conversion codes; unsigned long integer rather than unsigned integer when preceding the **o**, **u**, or **x** conversion codes.
    - h** Short integer rather than an integer when preceding the **d**, **i**, and **n** conversion codes; unsigned short integer (half integer) rather than an unsigned integer when preceding the **o**, **u**, or **x** conversion codes.
  - A conversion code that specifies the type of conversion to be applied.

The conversion specification takes the form:

```
%[*][width][size]convcode
```

*Format (Continued)*

The results from the conversion are placed in the memory location designated by the *Pointer* parameter unless you specify assignment suppression with an \* (asterisk). Assignment suppression provides a way to describe an input field to be skipped. The input field is a string of nonwhite space characters. It extends to the next inappropriate character or until the field width, if specified, is exhausted.

The conversion code indicates how to interpret the input field. The corresponding *Pointer* parameter must be a restricted type. Do not specify the *Pointer* parameter for a suppressed field. You can use the following conversion codes:

- %** Accepts a single % (percent sign) input at this point; no assignment or conversion is done. The complete conversion specification should be %% (two percent signs).
- d** Accepts an optionally signed decimal integer with the same format as that expected for the subject sequence of the **strtol** subroutine with a value of **10** for the *base* parameter. If no size modifier is specified, the *Pointer* parameter should be a pointer to an integer.
- i** Accepts an optionally signed integer with the same format as that expected for the subject sequence of the **strtol** subroutine with a value of **0** for the *base* parameter. If no size modifier is specified, the *Pointer* parameter should be a pointer to an integer.
- u** Accepts an optionally signed decimal integer with the same format as that expected for the subject sequence of the **strtoul** subroutine with a value of **10** for the *base* parameter. If no size modifier is specified, the *Pointer* parameter should be a pointer to an unsigned integer.
- o** Accepts an optionally signed octal integer with the same format as that expected for the subject sequence of the **strtoul** subroutine with a value of **8** for the *base* parameter. If no size modifier is specified, the *Pointer* parameter should be a pointer to an unsigned integer.
- x** Accepts an optionally signed hexadecimal integer with the same format as that expected for the subject sequence of the **strtoul** subroutine with a value of **16** for the *base* parameter. If no size modifier is specified, the *Pointer* parameter should be a pointer to an integer.

**e, f, or g**

Accepts an optionally signed floating-point number with the same format as that expected for the subject sequence of the **strtod** subroutine. The next field is converted accordingly and stored through the corresponding parameter; if no size modifier is specified, this parameter should be a pointer to a float. The input format for floating-point numbers is a string of digits, with some optional characteristics:

- It can be a signed value.
- It can be an exponential value, containing a decimal rational number followed by an exponent field, which consists of an **E** or an **e** followed by an (optionally signed) integer.
- It can be one of the special values **INF**, **NaNQ**, or **NaNS**. This value is translated into the IEEE-754 value for infinity, quiet **NaN**, or signaling **NaN**, respectively.

- p** Matches an unsigned hexadecimal integer, the same as the **%p** conversion of the **printf** subroutine. The corresponding parameter is a pointer to a void pointer. If the input item is a value converted earlier during the same program execution, the resulting pointer compares equal to that value; otherwise, the results of the **%p** conversion are unpredictable.
- n** Consumes no input. The corresponding parameter is a pointer to an integer into which the **scanf**, **fscanf**, **sscanf**, or **wscanf** subroutine writes the number of characters (including wide characters) read from the input stream. The assignment count returned at the completion of this function is not incremented.
- s** Accepts a sequence of nonwhite space characters (**scanf**, **fscanf**, and **sscanf** subroutines). The **wscanf** subroutine accepts a sequence of nonwhite-space wide-character codes; this sequence is converted to a sequence of characters in the same manner as the **wcstombs** subroutine. The *Pointer* parameter should be a pointer to the initial byte of a **char**, signed **char**, or unsigned **char** array large enough to hold the sequence and a terminating null-character code, which is automatically added.
- S** Accepts a sequence of nonwhite space characters (**scanf**, **fscanf**, and **sscanf** subroutines). This sequence is converted to a sequence of wide-character codes in the same manner as the **mbstowcs** subroutine. The **wscanf** subroutine accepts a sequence of nonwhite-space wide character codes. The *Pointer* parameter should be a pointer to the initial wide character code of an

array large enough to accept the sequence and a terminating null wide character code, which is automatically added. If the field width is specified, it denotes the maximum number of characters to accept.

- c** Accepts a sequence of bytes of the number specified by the field width (**scanf**, **fscanf** and **sscanf** subroutines); if no field width is specified, 1 is the default. The **wsscanf** subroutine accepts a sequence of wide-character codes of the number specified by the field width; if no field width is specified, 1 is the default. The sequence is converted to a sequence of characters in the same manner as the **wcstombs** subroutine. The *Pointer* parameter should be a pointer to the initial bytes of an array large enough to hold the sequence; no null byte is added. The normal skip over white space does not occur.
- C** Accepts a sequence of characters of the number specified by the field width (**scanf**, **fscanf**, and **sscanf** subroutines); if no field width is specified, 1 is the default. The sequence is converted to a sequence of wide character codes in the same manner as the **mbstowcs** subroutine. The **wsscanf** subroutine accepts a sequence of wide-character codes of the number specified by the field width; if no field width is specified, 1 is the default. The *Pointer* parameter should be a pointer to the initial wide character code of an array large enough to hold the sequence; no null wide-character code is added.

#### [*scanset*]

Accepts a nonempty sequence of bytes from a set of expected bytes specified by the *scanset* variable (**scanf**, **fscanf**, and **sscanf** subroutines). The **wsscanf** subroutine accepts a nonempty sequence of wide-character codes from a set of expected wide-character codes specified by the *scanset* variable. The sequence is converted to a sequence of characters in the same manner as the **wcstombs** subroutine. The *Pointer* parameter should be a pointer to the initial character of a **char**, **signed char**, or **unsigned char** array large enough to hold the sequence and a terminating null byte, which is automatically added. In the **scanf**, **fscanf**, and **sscanf** subroutines, the conversion specification includes all subsequent bytes in the string specified by the *Format* parameter, up to and including the ] (right bracket). The bytes between the brackets comprise the *scanset* variable, unless the byte after the [ (left bracket) is a ^ (circumflex). In this case, the *scanset* variable contains all bytes that do not appear in the scanlist between the ^ (circumflex) and the ] (right bracket). In the **wsscanf** subroutine, the characters between the brackets are first converted to wide character codes in the same manner as the **mbtowc** subroutine. These wide character codes are then used as described above in place of the bytes in the scanlist. If the conversion specification begins with [] or [^], the right bracket is included in the scanlist and the next right bracket is the matching right bracket that ends the conversion specification. You can also:

- Represent a range of characters by the construct *First-Last*. Thus, you can express [0123456789] as [0-9]. The *First* parameter must be lexically less than or equal to the *Last* parameter or else the - (dash) stands for itself. The - also stands for itself whenever it is the first or the last character in the *scanset* variable.
- Include the ] (right bracket) as an element of the *scanset* variable if it is the first character of the *scanset*. In this case it is not interpreted as the bracket that closes the *scanset* variable. If the *scanset* variable is an exclusive *scanset* variable, the ] is preceded by the ^ (circumflex) to make the ] an element of the *scanset*. The corresponding *Pointer* parameter should point to a character array large enough to hold the data field and that ends with a null character (\0). The \0 is added automatically.

A **scanf** conversion ends at the end-of-file (EOF character), the end of the control string, or when an input character conflicts with the control string. If it ends with an input character conflict, the conflicting character is not read from the input stream.

Unless a match in the control string exists, trailing white space (including a new-line character) is not read.

The success of literal matches and suppressed assignments is not directly determinable.

The National Language Support (NLS) extensions to the **scanf** subroutines can handle a format string that enables the system to process elements of the argument list in variable order. The normal conversion character % is replaced by %n\$, where *n* is a decimal number. Conversions are then applied to the specified argument (that is, the *n*th argument), rather than to the next unused argument.

The first successful run of the **fgetc**, **fgets**, **fread**, **getc**, **getchar**, **gets**, **scanf**, or **fscanf** subroutine using a stream that returns data not supplied by a prior call to the **ungetc** (“ungetc or ungetwc Subroutine” on page 422) subroutine marks the `st_atime` field for update.

## Return Values

These subroutines return the number of successfully matched and assigned input items. This number can be 0 if an early conflict existed between an input character and the control string. If the input ends before the first conflict or conversion, only EOF is returned. If a read error occurs, the error indicator for the stream is set, EOF is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **scanf**, **fscanf**, **sscanf**, and **wscanf** subroutines are unsuccessful if either the file specified by the *Stream*, *String*, or *wcs* parameter is unbuffered or data needs to be read into the file’s buffer and one or more of the following conditions is true:

- EAGAIN** The **O\_NONBLOCK** flag is set for the file descriptor underlying the file specified by the *Stream*, *String*, or *wcs* parameter, and the process would be delayed in the **scanf**, **fscanf**, **sscanf**, or **wscanf** operation.
- EBADF** The file descriptor underlying the file specified by the *Stream*, *String*, or *wcs* parameter is not a valid file descriptor open for reading.
- EINTR** The read operation was terminated due to receipt of a signal, and either no data was transferred or a partial transfer was not reported.

**Note:** Depending upon which library routine the application binds to, this subroutine may return **EINTR**. Refer to the **signal** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine regarding **SA\_RESTART**.

- EIO** The process is a member of a background process group attempting to perform a read from its controlling terminal, and either the process is ignoring or blocking the **SIGTTIN** signal or the process group has no parent process.
- EINVAL** The subroutine received insufficient arguments for the *Format* parameter.
- EILSEQ** A character sequence that is not valid was detected, or a wide-character code does not correspond to a valid character.
- ENOMEM** Insufficient storage space is available.

## Related Information

The **atof**, **atoff**, **strtod**, or **strtof** subroutine, **fread** subroutine, **getc**, **fgetc**, **getchar**, or **getw** subroutine, **gets** or **fgets** subroutine, **getwc**, **fgetwc**, or **getwchar** subroutine, **mbstowcs** subroutine, **mbtowc** subroutine, **printf**, **fprintf**, **sprintf**, **wsprintf**, **vprintf**, **vfprintf**, **vsprintf**, or **vwsprintf** subroutine, **setlocale** (“setlocale Subroutine” on page 136) subroutine, **strtol**, **strtoul**, **atol**, or **atoi** (“strtol, strtoul, strtoll, strtoull, or atoi Subroutine” on page 299) subroutine, **ungetc** (“ungetc or ungetwc Subroutine” on page 422) subroutine, **wcstombs** (“wcstombs Subroutine” on page 461) subroutine.

Input and Output Handling Programmer’s Overview, in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview for Programming in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## sched\_yield Subroutine

### Purpose

Yield processor.

### Library

Standard Library (**libc.a**)

### Syntax

```
#include <sched.h>
int sched_yield (void) ;
```

### Description

The **sched\_yield** function forces the running thread to relinquish the processor until it again becomes the head of its thread list. It takes no arguments.

### Return Values

The **sched\_yield** function returns 0 if it completes successfully, or it returns a value of -1 and sets `errno` to indicate the error.

---

## select Subroutine

### Purpose

Checks the I/O status of multiple file descriptors and message queues.

### Library

Standard C Library (**libc.a**)

### Syntax

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

int select (Nfdsmsgs, ReadList, WriteList, ExceptList, TimeOut)
int Nfdsmsgs;
struct sellist * ReadList, *WriteList, *ExceptList;
struct timeval * TimeOut;
```

### Description

The **select** subroutine checks the specified file descriptors and message queues to see if they are ready for reading (receiving) or writing (sending), or if they have an exceptional condition pending.

When selecting on an unconnected stream socket, **select** returns when the connection is made. If selecting on a connected stream socket, then the ready message indicates that data can be sent or received. Files descriptors of regular files always select true for read, write, and exception conditions. For more information on sockets, refer to "Understanding Socket Connections" and the related "Checking for Pending Connections Example Program" dealing with pending connections in *AIX 5L Version 5.2 Communications Programming Concepts*.

The **select** subroutine is also supported for compatibility with previous releases of this operating system and with BSD systems.

**Note:** If selecting on a non-blocking socket for both read and write events and if the destination host is unreachable, **select** could show a different behavior due to timing constraints. Refer to the Examples section of this document for further information..

## Parameters

*Nfdsmsgs* Specifies the number of file descriptors and the number of message queues to check. The low-order 16 bits give the length of a bit mask that specifies which file descriptors to check; the high-order 16 bits give the size of an array that contains message queue identifiers. If either half of the *Nfdsmsgs* parameter is equal to a value of 0, the corresponding bit mask or array is assumed not to be present.

*TimeOut* Specifies either a null pointer or a pointer to a **timeval** structure that specifies the maximum length of time to wait for at least one of the selection criteria to be met. The **timeval** structure is defined in the `/usr/include/sys/time.h` file and it contains the following members:

```
struct timeval {
    int tv_sec;          /* seconds      */
    int tv_usec;        /* microseconds */
};
```

The number of microseconds specified in *TimeOut.tv\_usec*, a value from 0 to 999999, is set to one millisecond if the process does not have root user authority and the value is less than one millisecond.

If the *TimeOut* parameter is a null pointer, the **select** subroutine waits indefinitely, until at least one of the selection criteria is met. If the *TimeOut* parameter points to a **timeval** structure that contains zeros, the file and message queue status is polled, and the **select** subroutine returns immediately.

*ReadList, WriteList,  
ExceptList*

Specify what to check for reading, writing, and exceptions, respectively. Together, they specify the selection criteria. Each of these parameters points to a **sellist** structure, which can specify both file descriptors and message queues. Your program must define the **sellist** structure in the following form:

```
struct sellist
{
int fdsmask[F];          /* file descriptor bit mask */
int msgids[M];          /* message queue identifiers */
};
```

The `fdsmask` array is treated as a bit string in which each bit corresponds to a file descriptor. File descriptor  $n$  is represented by the bit  $(1 \ll (n \bmod \text{bits}))$  in the array element `fdsmask[n / BITS(int)]`. (The **BITS** macro is defined in the **values.h** file.) Each bit that is set to 1 indicates that the status of the corresponding file descriptor is to be checked.

**Note:** The low-order 16 bits of the `Nfdsmgs` parameter specify the number of *bits* (not elements) in the `fdsmask` array that make up the file descriptor mask. If only part of the last int is included in the mask, the appropriate number of low-order bits are used, and the remaining high-order bits are ignored. If you set the low-order 16 bits of the `Nfdsmgs` parameter to 0, you must *not* define an `fdsmask` array in the **sellist** structure.

Each int of the `msgids` array specifies a message queue identifier whose status is to be checked. Elements with a value of -1 are ignored. The high-order 16 bits of the `Nfdsmgs` parameter specify the number of elements in the `msgids` array. If you set the high-order 16 bits of the `Nfdsmgs` parameter to 0, you must *not* define a `msgids` array in the **sellist** structure.

**Note:** The arrays specified by the *ReadList*, *WriteList*, and *ExceptList* parameters are the same size because each of these parameters points to the same **sellist** structure type. However, you need not specify the same number of file descriptors or message queues in each. Set the file descriptor bits that are not of interest to 0, and set the extra elements of the `msgids` array to -1.

You can use the **SELLIST** macro defined in the **sys/select.h** file to define the **sellist** structure. The format of this macro is:

```
SELLIST(f, m) declarator . . . ;
```

where *f* specifies the size of the `fdsmask` array, *m* specifies the size of the `msgids` array, and each *declarator* is the name of a variable to be declared as having this type.

## Return Values

Upon successful completion, the **select** subroutine returns a value that indicates the total number of file descriptors and message queues that satisfy the selection criteria. The `fdsmask` bit masks are modified so that bits set to 1 indicate file descriptors that meet the criteria. The `msgids` arrays are altered so that message queue identifiers that do not meet the criteria are replaced with a value of -1.

The return value is similar to the `Nfdsmgs` parameter in that the low-order 16 bits give the number of file descriptors, and the high-order 16 bits give the number of message queue identifiers. These values indicate the sum total that meet each of the read, write, and exception criteria. Therefore, the same file descriptor or message queue can be counted up to three times. You can use the **NFDS** and **NMSGs** macros found in the **sys/select.h** file to separate out these two values from the return value. For example, if *rc* contains the value returned from the **select** subroutine, **NFDS**(*rc*) is the number of files selected, and **NMSGs**(*rc*) is the number of message queues selected.

If the time limit specified by the *TimeOut* parameter expires, the **select** subroutine returns a value of 0.

If a connection-based socket is specified in the *Readlist* parameter and the connection disconnects, the **select** subroutine returns successfully, but the **recv** subroutine on the socket will return a value of 0 to indicate the socket connection has been closed.

For nonblocking connection-based sockets, both successful and unsuccessful connections will cause the **select** subroutine to return successfully without any error.

When the connection completes successfully the socket becomes writable, and if the connection encounters an error the socket becomes both readable and writable.

When using the **select** subroutine, you can not check any pending errors on the socket. You need to call the **getsockopt** subroutine with **SOL\_SOCKET** and **SOL\_ERROR** to check for a pending error.

If the **select** subroutine is unsuccessful, it returns a value of -1 and sets the global variable **errno** to indicate the error. In this case, the contents of the structures pointed to by the *ReadList*, *WriteList*, and *ExceptList* parameters are unpredictable.

## Error Codes

The **select** subroutine is unsuccessful if one of the following are true:

<b>EBADF</b>	An invalid file descriptor or message queue identifier was specified.
<b>EAGAIN</b>	Allocation of internal data structures was unsuccessful.
<b>EINTR</b>	A signal was caught during the <b>select</b> subroutine and the signal handler was installed with an indication that subroutines are not to be restarted.
<b>EINVAL</b>	One of the parameters to the <b>select</b> subroutine contained a value that is not valid.
<b>EFAULT</b>	The <i>ReadList</i> , <i>WriteList</i> , <i>ExceptList</i> , or <i>Timeout</i> parameter points to a location outside of the address space of the process.

## Examples

The following is an example of the behavior of the **select** subroutine called on a non-blocking socket, when trying to connect to a host that is unreachable:

```
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <netinet/tcp.h>
#include <fcntl.h>
#include <sys/time.h>
#include <errno.h>
#include <stdio.h>

int main()
{
    int sockfd, cnt, i = 1;
    struct sockaddr_in serv_addr;

    bzero((char *)&serv_addr, sizeof (serv_addr));
    serv_addr.sin_family = AF_INET;
    serv_addr.sin_addr.s_addr = inet_addr("172.16.55.25");
    serv_addr.sin_port = htons(102);

    if ((sockfd = socket(AF_INET, SOCK_STREAM, 0)) < 0)
        exit(1);
    if (fcntl(sockfd, F_SETFL, FNONBLOCK) < 0)
        exit(1);
    if (connect(sockfd, (struct sockaddr *)&serv_addr, sizeof
        (serv_addr)) < 0 && errno != EINPROGRESS)
        exit(1);
    for (cnt=0; cnt<2; cnt++) {
        fd_set readfds, writefds;

        FD_ZERO(&readfds);
        FD_SET(sockfd, &readfds);
        FD_ZERO(&writefds);
        FD_SET(sockfd, &writefds);
```

```

        if (select(sockfd + 1, &readfds, &writefds, NULL,
                NULL) < 0)
            exit(1);
        printf("Iteration %d =====\n", i);
        printf("FD_ISSET(sockfd, &readfds) == %d\n",
                FD_ISSET(sockfd, &readfds));
        printf("FD_ISSET(sockfd, &writefds) == %d\n",
                FD_ISSET(sockfd, &writefds));
        i++;
    }
    return 0;
}

```

Here is the output of the above program :

```

Iteration 1 =====
FD_ISSET(sockfd, &readfds) == 0
FD_ISSET(sockfd, &writefds) == 1
Iteration 2 =====
FD_ISSET(sockfd, &readfds) == 1
FD_ISSET(sockfd, &writefds) == 1

```

In the first iteration, **select** notifies the write event only. In the second iteration, **select** notifies both the read and write events.

## Notes

**FD\_SETSIZE** is the #define variable that defines how many file descriptors the various FD macros will use. The default value for **FD\_SETSIZE** will vary, depending on the version of AIX. As the number of open files supported has increased, the default value of **FD\_SETSIZE** has increased.

In AIX Version 4.3.1, the size increased to 32767 open file descriptors (from 2000 in prior releases). In AIX 5L Version 5.2.0, the size increased to 65534 open file descriptors. This value can not be set greater than **OPEN\_MAX**, which also varies from one AIX Version to another.

For more information, refer to the `/usr/include/sys/time.h` file.

The user may override **FD\_SETSIZE** to select a smaller value before including the system header files. This is desirable for performance reasons, because of the overhead in **FD\_ZERO** to zero 65534 bits.

## Performance Issues and Recommended Coding Practices

The **select** subroutine can be a very compute intensive system call, depending on the number of open file descriptors used and the lengths of the bit maps used. Do not follow the examples shown in many text books. Most were written when the number of open files supported was small, and thus the bit maps were short. You should avoid the following (where **select** is being passed **FD\_SETSIZE** as the number of FDs to process):

```
select(FD_SETSIZE, ....)
```

Performance will be poor if the program uses **FD\_ZERO** and the default **FD\_SETSIZE**. **FD\_ZERO** should not be used in any loops or before each **select** call. However, using it one time to zero the bit string will not cause problems. If you plan to use this simple programming method, you should override **FD\_SETSIZE** to define a smaller number of FDs. For example, if your process will only open two FDs that you will be selecting on, and there will never be more than a few hundred other FDs open in the process, you should lower **FD\_SETSIZE** to approximately 1024.

Do not pass **FD\_SETSIZE** as the first parameter to **select**. This specifies the maximum number of file descriptors the system should check for. The program should keep track of the highest FD that has been assigned or use the **getdtablesize** subroutine to determine this value. This saves passing excessively long bit maps in and out of the kernel and reduces the number of FDs that **select** must check.

Use the **poll** system call instead of **select**. The **poll** system call has the same functionality as **select**, but it uses a list of FDs instead of a bit map. Thus, if you are only selecting on a single FD, you would only pass one FD to **poll**. With **select**, you have to pass a bit map that is as long as the FD number assigned for that FD. If AIX assigned FD 4000, for example, you would have to pass a bit map 4001 bits long.

## Related Information

The **poll** subroutine.

The Input and Output Handling Programmer's Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## semctl Subroutine

### Purpose

Controls semaphore operations.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int semctl (SemaphoreID, SemaphoreNumber, Command, arg)
```

OR

```
int semctl (SemaphoreID, SemaphoreNumber, Command)
```

```
int SemaphoreID;  
int SemaphoreNumber;  
int Command;  
union semun {  
    int val;  
    struct semid_ds *buf;  
    unsigned short *array;  
} arg;
```

If the fourth argument is required for the operation requested, it must be of type union semun and explicitly declared as shown above.

### Description

The **semctl** subroutine performs a variety of semaphore control operations as specified by the *Command* parameter.

The following limits apply to semaphores:

- Maximum number of semaphore IDs is 4096 for operating system releases before AIX 4.3.2 and 131072 for AIX 4.3.2 and following.
- Maximum number of semaphores per ID is 65,535.
- Maximum number of operations per call by the **semop** (“semop Subroutine” on page 125) subroutine is 1024.
- Maximum number of undo entries per procedure is 1024.
- Maximum semaphore value is 32,767.
- Maximum adjust-on-exit value is 16,384.

## Parameters

### *SemaphoreID*

Specifies the semaphore identifier.

### *SemaphoreNumber*

Specifies the semaphore number.

*arg.val* Specifies the value for the semaphore for the **SETVAL** command.

### *arg.buf*

Specifies the buffer for status information for the **IPC\_STAT** and **IPC\_SET** commands.

### *arg.array*

Specifies the values for all the semaphores in a set for the **GETALL** and **SETALL** commands.

### *Command*

Specifies semaphore control operations.

The following *Command* parameter values are executed with respect to the semaphore specified by the *SemaphoreID* and *SemaphoreNumber* parameters. These operations get and set the values of a **sem** structure, which is defined in the **sys/sem.h** file.

#### **GETVAL**

Returns the **semval** value, if the current process has read permission.

#### **SETVAL**

Sets the **semval** value to the value specified by the *arg.val* parameter, if the current process has write permission. When this *Command* parameter is successfully executed, the **semadj** value corresponding to the specified semaphore is cleared in all processes.

#### **GETPID**

Returns the value of the **sempid** field, if the current process has read permission.

#### **GETNCNT**

Returns the value of the **semncnt** field, if the current process has read permission.

#### **GETZCNT**

Returns the value of the **semzcnt** field, if the current process has read permission.

The following *Command* parameter values return and set every **semval** value in the set of semaphores. These operations get and set the values of a **sem** structure, which is defined in the **sys/sem.h** file.

#### **GETALL**

Stores **semvals** values into the array pointed to by the *arg.array* parameter, if the current process has read permission.

#### **SETALL**

Sets **semvals** values according to the array pointed to by the *arg.array* parameter, if the current process has write permission. When this *Command* parameter is successfully executed, the **semadj** value corresponding to each specified semaphore is cleared in all processes.

The following *Commands* parameter values get and set the values of a **semid\_ds** structure, defined in the **sys/sem.h** file. These operations get and set the values of a **sem** structure, which is defined in the **sys/sem.h** file.

#### **IPC\_STAT**

Obtains status information about the semaphore identified by the *SemaphoreID* parameter. This information is stored in the area pointed to by the *arg.buf* parameter.

#### **IPC\_SET**

Sets the owning user and group IDs, and the access permissions for the set of

semaphores associated with the *SemaphoreID* parameter. The **IPC\_SET** operation uses as input the values found in the *arg.buf* parameter structure.

**IPC\_SET** sets the following fields:

sem_perm.uid	User ID of the owner
sem_perm.gid	Group ID of the owner
sem_perm.mode	Permission bits only
sem_perm.cuid	Creator's user ID

**IPC\_SET** can only be executed by a process that has root user authority or an effective user ID equal to the value of the *sem\_perm.uid* or *sem\_perm.cuid* field in the data structure associated with the *SemaphoreID* parameter.

### IPC\_RMID

Removes the semaphore identifier specified by the *SemaphoreID* parameter from the system and destroys the set of semaphores and data structures associated with it. This *Command* parameter can only be executed by a process that has root user authority or an effective user ID equal to the value of the *sem\_perm.uid* or *sem\_perm.cuid* field in the data structure associated with the *SemaphoreID* parameter.

## Return Values

Upon successful completion, the value returned depends on the *Command* parameter as follows:

Command	Return Value
<b>GETVAL</b>	Returns the value of the <i>semval</i> field.
<b>GETPID</b>	Returns the value of the <i>sempid</i> field.
<b>GETNCNT</b>	Returns the value of the <i>semncnt</i> field.
<b>GETZCNT</b>	Returns the value of the <i>semzcnt</i> field.
<b>All Others</b>	Return a value of 0.

If the **semctl** subroutine is unsuccessful, a value of -1 is returned and the global variable **errno** is set to indicate the error.

## Error Codes

The **semctl** subroutine is unsuccessful if any of the following is true:

<b>EINVAL</b>	The <i>SemaphoreID</i> parameter is not a valid semaphore identifier.
<b>EINVAL</b>	The <i>SemaphoreNumber</i> parameter is less than 0 or greater than or equal to the <b>sem_nsems</b> value.
<b>EINVAL</b>	The <i>Command</i> parameter is not a valid command.
<b>EACCES</b>	The calling process is denied permission for the specified operation.
<b>ERANGE</b>	The <i>Command</i> parameter is equal to the <b>SETVAL</b> or <b>SETALL</b> value and the value to which <b>semval</b> value is to be set is greater than the system-imposed maximum.
<b>EPERM</b>	The <i>Command</i> parameter is equal to the <b>IPC_RMID</b> or <b>IPC_SET</b> value and the calling process does not have root user authority or an effective user ID equal to the value of the <i>sem_perm.uid</i> or <i>sem_perm.cuid</i> field in the data structure associated with the <i>SemaphoreID</i> parameter.
<b>EFAULT</b>	The <i>arg.buf</i> or <i>arg.array</i> parameter points outside of the allocated address space of the process.
<b>ENOMEM</b>	The system does not have enough memory to complete the subroutine.

## Related Information

The **semget** (“semget Subroutine” on page 123) subroutine, **semop** (“semop Subroutine” on page 125) subroutine.

---

## semget Subroutine

### Purpose

Gets a set of semaphores.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int semget (Key, NumberOfSemaphores, SemaphoreFlag)
```

```
key_t Key;
```

```
int NumberOfSemaphores, SemaphoreFlag;
```

### Description

The **semget** subroutine returns the semaphore identifier associated with the *Key* parameter value.

The **semget** subroutine creates a data structure for the semaphore ID and an array containing the *NumberOfSemaphores* parameter semaphores if one of the following conditions is true:

- The *Key* parameter is equal to the **IPC\_PRIVATE** operation.
- The *Key* parameter does not already have a semaphore identifier associated with it, and the **IPC\_CREAT** value is set.

Upon creation, the data structure associated with the new semaphore identifier is initialized as follows:

- The *sem\_perm.cuid* and *sem\_perm.uid* fields are set equal to the effective user ID of the calling process.
- The *sem\_perm.cgid* and *sem\_perm.gid* fields are set equal to the effective group ID of the calling process.
- The low-order 9 bits of the *sem\_perm.mode* field are set equal to the low-order 9 bits of the *SemaphoreFlag* parameter.
- The *sem\_nsems* field is set equal to the value of the *NumberOfSemaphores* parameter.
- The *sem\_otime* field is set equal to 0 and the *sem\_ctime* field is set equal to the current time.

The data structure associated with each semaphore in the set is not initialized. The **semctl** (“semctl Subroutine” on page 120) subroutine (with the *Command* parameter values **SETVAL** or **SETALL**) can be used to initialize each semaphore.

If the *Key* parameter value is not **IPC\_PRIVATE**, the **IPC\_EXCL** value is not set, and a semaphore identifier already exists for the specified *Key* parameter, the value of the *NumberOfSemaphores* parameter specifies the number of semaphores that the current process needs.

If the *NumberOfSemaphores* parameter has a value of 0, any number of semaphores is acceptable. If the *NumberOfSemaphores* parameter is not 0, the **semget** subroutine is unsuccessful if the set contains fewer than the value of the *NumberOfSemaphores* parameter.

The following limits apply to semaphores:

- Maximum number of semaphore IDs is 4096 for operating system releases before AIX 4.3.2 and 131072 for AIX 4.3.2 and following.
- Maximum number of semaphores per ID is 65,535.
- Maximum number of operations per call by the **semop** subroutine is 1024.
- Maximum number of undo entries per procedure is 1024.

- Maximum semaphore value is 32,767.
- Maximum adjust-on-exit value is 16,384.

## Parameters

<i>Key</i>	Specifies either the <b>IPC_PRIVATE</b> value or an IPC key constructed by the <b>ftok</b> subroutine (or a similar algorithm).
<i>NumberOfSemaphores</i>	Specifies the number of semaphores in the set.
<i>SemaphoreFlag</i>	Constructed by logically ORing one or more of the following values: <ul style="list-style-type: none"> <li><b>IPC_CREAT</b> Creates the data structure if it does not already exist.</li> <li><b>IPC_EXCL</b> Causes the <b>semget</b> subroutine to fail if the <b>IPC_CREAT</b> value is also set and the data structure already exists.</li> <li><b>S_IRUSR</b> Permits the process that owns the data structure to read it.</li> <li><b>S_IWUSR</b> Permits the process that owns the data structure to modify it.</li> <li><b>S_IRGRP</b> Permits the group associated with the data structure to read it.</li> <li><b>S_IWGRP</b> Permits the group associated with the data structure to modify it.</li> <li><b>S_IROTH</b> Permits others to read the data structure.</li> <li><b>S_IWOTH</b> Permits others to modify the data structure.</li> </ul>
	Values that begin with the <b>S_I</b> prefix are defined in the <b>sys/mode.h</b> file and are a subset of the access permissions that apply to files.

## Return Values

Upon successful completion, the **semget** subroutine returns a semaphore identifier. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **semget** subroutine is unsuccessful if one or more of the following conditions is true:

<b>EACCES</b>	A semaphore identifier exists for the <i>Key</i> parameter but operation permission, as specified by the low-order 9 bits of the <i>SemaphoreFlag</i> parameter, is not granted.
<b>EINVAL</b>	A semaphore identifier does not exist and the <i>NumberOfSemaphores</i> parameter is less than or equal to a value of 0, or greater than the system-imposed value.
<b>EINVAL</b>	A semaphore identifier exists for the <i>Key</i> parameter, but the number of semaphores in the set associated with it is less than the value of the <i>NumberOfSemaphores</i> parameter and the <i>NumberOfSemaphores</i> parameter is not equal to 0.
<b>ENOENT</b>	A semaphore identifier does not exist for the <i>Key</i> parameter and the <b>IPC_CREAT</b> value is not set.
<b>ENOSPC</b>	Creating a semaphore identifier would exceed the maximum number of identifiers allowed systemwide.
<b>EEXIST</b>	A semaphore identifier exists for the <i>Key</i> parameter, but both the <b>IPC_CREAT</b> and <b>IPC_EXCL</b> values are set.
<b>ENOMEM</b>	There is not enough memory to complete the operation.

## Related Information

The **ftok** subroutine, **semctl** (“semctl Subroutine” on page 120) subroutine, **semop** (“semop Subroutine”) subroutine.

The **mode.h** file.

---

## semop Subroutine

### Purpose

Performs semaphore operations.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int semop (SemaphoreID, SemaphoreOperations, NumberOfSemaphoreOperations)
int SemaphoreID;
struct sembuf * SemaphoreOperations;
size_t NumberOfSemaphoreOperations;
```

### Description

The **semop** subroutine performs operations on the set of semaphores associated with the semaphore identifier specified by the *SemaphoreID* parameter. The **sembuf** structure is defined in the **usr/include/sys/sem.h** file.

Each **sembuf** structure specified by the *SemaphoreOperations* parameter includes the following fields:

sem_num	Semaphore number
sem_op	Semaphore operation
sem_flg	Operation flags

Each semaphore operation specified by the *sem\_op* field is performed on the semaphore specified by the *SemaphoreID* parameter and the *sem\_num* field. The *sem\_op* field specifies one of three semaphore operations.

1. If the *sem\_op* field is a negative integer and the calling process has permission to alter, one of the following conditions occurs:
  - If the **semval** variable (see the **usr/include/sys/sem.h** file) is greater than or equal to the absolute value of the *sem\_op* field, the absolute value of the *sem\_op* field is subtracted from the **semval** variable. In addition, if the **SEM\_UNDO** flag is set in the *sem\_flg* field, the absolute value of the *sem\_op* field is added to the **semadj** value of the calling process for the specified semaphore.
  - If the **semval** variable is less than the absolute value of the *sem\_op* field and the **IPC\_NOWAIT** value is set in the *sem\_flg* field, the **semop** subroutine returns immediately.
  - If the **semval** variable is less than the absolute value of the *sem\_op* field and the **IPC\_NOWAIT** value is not set in the *sem\_flg* field, the **semop** subroutine increments the *semncnt* field associated with the specified semaphore and suspends the calling process until one of the following conditions occurs:
    - The value of the **semval** variable becomes greater than or equal to the absolute value of the *sem\_op* field. The value of the *semncnt* field associated with the specified semaphore is then decremented, and the absolute value of the *sem\_op* field is subtracted from the **semval** variable.

In addition, if the **SEM\_UNDO** flag is set in the `sem_flg` field, the absolute value of the `sem_op` field is added to the **semadj** value of the calling process for the specified semaphore.

- The *SemaphoreID* parameter for which the calling process is awaiting action is removed from the system. When this occurs, the **errno** global variable is set to the **EIDRM** flag and a value of -1 is returned.
  - The calling process received a signal that is to be caught. When this occurs, the **semop** subroutine decrements the value of the `semncnt` field associated with the specified semaphore. When the `semzcnt` field is decremented, the calling process resumes as prescribed by the **sigaction** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine.
2. If the `sem_op` field is a positive integer and the calling process has alter permission, the value of the `sem_op` field is added to the **semval** variable. In addition, if the **SEM\_UNDO** flag is set in the `sem_flg` field, the value of the `sem_op` field is subtracted from the calling process’s **semadj** value for the specified semaphore.
  3. If the value of the `sem_op` field is 0 and the calling process has read permission, one of the following occurs:
    - If the **semval** variable is 0, the **semop** subroutine returns immediately.
    - If the **semval** variable is not equal to 0 and **IPC\_NOWAIT** value is set in the `sem_flg` field, the **semop** subroutine returns immediately.
    - If the **semval** variable is not equal to 0 and the **IPC\_NOWAIT** value is set in the `sem_flg` field, the **semop** subroutine increments the `semzcnt` field associated with the specified semaphore and suspends execution of the calling process until one of the following occurs:
      - The value of the **semval** variable becomes 0. When this occurs, the value of the `semzcnt` field associated with the specified semaphore is decremented.
      - The *SemaphoreID* parameter for which the calling process is awaiting action is removed from the system. If this occurs, the **errno** global variable is set to the **EIDRM** error code and a value of -1 is returned.
      - The calling process received a signal that is to be caught. When this occurs, the **semop** subroutine decrements the value of the `semzcnt` field associated with the specified semaphore. When the `semzcnt` field is decremented, the calling process resumes execution as prescribed by the **sigaction** subroutine.

The following limits apply to semaphores:

- Maximum number of semaphore IDs is 4096 for operating system releases before AIX 4.3.2 and 131072 for AIX 4.3.2 and following.
- Maximum number of semaphores per ID is 65,535.
- Maximum number of operations per call by the **semop** subroutine is 1024.
- Maximum number of undo entries per procedure is 1024.
- Maximum capacity of a semaphore value is 32,767 bytes.
- Maximum adjust-on-exit value is 16,384 bytes.

## Parameters

<i>SemaphoreID</i>	Specifies the semaphore identifier.
<i>NumberOfSemaphoreOperations</i>	Specifies the number of structures in the array.
<i>SemaphoreOperations</i>	Points to an array of structures, each of which specifies a semaphore operation.

## Return Values

Upon successful completion, the **semop** subroutine returns a value of 0. Also, the *SemaphoreID* parameter value for each semaphore that is operated upon is set to the process ID of the calling process.

If the **semop** subroutine is unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error. If the **SEM\_ORDER** flag was set in the `sem_flg` field for the first semaphore operation in the *SemaphoreOperations* array, the **SEM\_ERR** value is set in the `sem_flg` field for the unsuccessful operation.

If the *SemaphoreID* parameter for which the calling process is awaiting action is removed from the system, the **errno** global variable is set to the **EIDRM** error code and a value of -1 is returned.

## Error Codes

The **semop** subroutine is unsuccessful if one or more of the following are true for any of the semaphore operations specified by the *SemaphoreOperations* parameter. If the operations were performed individually, the discussion of the **SEM\_ORDER** flag provides more information about error situations.

<b>EINVAL</b>	The <i>SemaphoreID</i> parameter is not a valid semaphore identifier.
<b>EFBIG</b>	The <code>sem_num</code> value is less than 0 or it is greater than or equal to the number of semaphores in the set associated with the <i>SemaphoreID</i> parameter.
<b>E2BIG</b>	The <i>NumberOfSemaphoreOperations</i> parameter is greater than the system-imposed maximum.
<b>EACCES</b>	The calling process is denied permission for the specified operation.
<b>EAGAIN</b>	The operation would result in suspension of the calling process, but the <b>IPC_NOWAIT</b> value is set in the <code>sem_flg</code> field.
<b>ENOSPC</b>	The limit on the number of individual processes requesting a <b>SEM_UNDO</b> flag would be exceeded.
<b>EINVAL</b>	The number of individual semaphores for which the calling process requests a <b>SEM_UNDO</b> flag would exceed the limit.
<b>ERANGE</b>	An operation would cause a <b>semval</b> value to overflow the system-imposed limit.
<b>ERANGE</b>	An operation would cause a <b>semadj</b> value to overflow the system-imposed limit.
<b>EFAULT</b>	The <i>SemaphoreOperations</i> parameter points outside of the address space of the process.
<b>EINTR</b>	A signal interrupted the <b>semop</b> subroutine.
<b>EIDRM</b>	The semaphore identifier <i>SemaphoreID</i> parameter has been removed from the system.

## Related Information

The **exec** subroutine, **exit** subroutine, **fork** subroutine, **semctl** (“semctl Subroutine” on page 120) subroutine, **semget** (“semget Subroutine” on page 123) subroutine, **sigaction** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine.

---

## setacldb or endacldb Subroutine

### Purpose

Opens and closes the SMIT ACL database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
int setacldb(Mode)
int Mode;
int endacldb;
```

### Description

These functions may be used to open and close access to the user SMIT ACL database. Programs that call the **getusraclattr** or **getgrpacclattr** subroutines should call the **setacldb** subroutine to open the database and the **endacldb** subroutine to close the database.

The **setacldb** subroutine opens the database in the specified mode, if it is not already open. The open count is increased by 1.

The **endacldb** subroutine decreases the open count by 1 and closes the database when this count goes to 0. Any uncommitted changed data is lost.

## Parameters

*Mode* Specifies the mode of the open. This parameter may contain one or more of the following values defined in the **usersec.h** file:

**S\_READ**

Specifies read access.

**S\_WRITE**

Specifies update access.

## Return Values

The **setacldb** and **endacldb** subroutines return a value of 0 to indicate success. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **setacldb** subroutine fails if the following is true:

**EACCES** Access permission is denied for the data request.

Both subroutines return errors from other subroutines.

## Security

Security Files Accessed: The calling process must have access to the SMIT ACL data.

Mode File **rw/etc/security/smitacl.user**

## Related Information

The **getgrpaclattr**, **nextgrpacl**, or **putgrpaclattr** subroutine, **getusraclattr**, **nextusracl**, or **putusraclattr** subroutine.

---

## setauthdb Subroutine

### Purpose

Restricts the search order for loadable authentication modules.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int setauthdb (new_db_name, old_db_name)
const char *new_db_name;
char *old_db_name;
```

## Description

The **setauthdb** library function controls which loadable authentication modules are examined by the security library functions. Specifying a value for the *new\_auth\_db* parameter will restrict the security library functions to the named simple or compound loadable authentication module. Future calls to security library functions will be restricted to the given loadable authentication module. The restriction may be removed by using a zero length string as the value of the *new\_auth\_db* parameter.

The current value may be obtained by providing a pointer to a character array which is long enough to hold the longest permissible stanza name in the **/usr/lib/security/methods.cfg** file (15 characters).

## Parameters

*new\_db\_name*

Pointer to the name of the new database module. The *new\_db\_name* parameter must reference a value module name contained in the **/usr/lib/security/methods.cfg** file, or one of the predefined values (*BUILTIN*, *compat*, or *files*). The empty string may be used to remove the restriction on which modules are used.

*old\_db\_name*

Pointer to where the name of the current module will be stored. A NULL value for the *old\_db\_name* parameter may be used if the current name of the database is not wanted.

## Return Values

0

The module search restriction has been successfully changed.

-1

The module search restriction could not be changed. The **errno** variable must be examined to determine the cause of the failure.

## Error Codes

**EINVAL**

The *new\_auth\_db* parameter is longer than the permissible length of a stanza in the **/usr/lib/security/methods.cfg** file (15 characters).

**ENOENT**

The *new\_auth\_db* does not reference a valid stanza in **/usr/lib/security/methods.cfg** or one of the predefined values.

## Related Information

getauthdb Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

---

## setbuf, setvbuf, setbuffer, or setlinebuf Subroutine

### Purpose

Assigns buffering to a stream.

### Library

Standard C Library (**libc.a**)

## Syntax

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

```
void setbuf ( Stream, Buffer)
```

```
FILE *Stream;
```

```
char *Buffer;
```

```
int setvbuf (Stream, Buffer, Mode, Size)
```

```
FILE *Stream;
```

```
char *Buffer;
```

```
int Mode;
```

```
size_t Size;
```

```
void setbuffer (Stream, Buffer, Size)
```

```
FILE *Stream;
```

```
char *Buffer;
```

```
size_t Size;
```

```
void setlinebuf (Stream)
```

```
FILE *Stream;
```

## Description

The **setbuf** subroutine causes the character array pointed to by the *Buffer* parameter to be used instead of an automatically allocated buffer. Use the **setbuf** subroutine after a stream has been opened, but before it is read or written.

If the *Buffer* parameter is a null character pointer, input/output is completely unbuffered.

A constant, **BUFSIZ**, defined in the **stdio.h** file, tells how large an array is needed:

```
char buf[BUFSIZ];
```

For the **setvbuf** subroutine, the *Mode* parameter determines how the *Stream* parameter is buffered:

**\_IOFBF** Causes input/output to be fully buffered.

**\_IOLBF** Causes output to be line-buffered. The buffer is flushed when a new line is written, the buffer is full, or input is requested.

**\_IONBF** Causes input/output to be completely unbuffered.

If the *Buffer* parameter is not a null character pointer, the array it points to is used for buffering. The *Size* parameter specifies the size of the array to be used as a buffer, but all of the *Size* parameter's bytes are not necessarily used for the buffer area. The constant **BUFSIZ** in the **stdio.h** file is one buffer size. If input/output is unbuffered, the subroutine ignores the *Buffer* and *Size* parameters. The **setbuffer** subroutine, an alternate form of the **setbuf** subroutine, is used after *Stream* has been opened, but before it is read or written. The character array *Buffer*, whose size is determined by the *Size* parameter, is used instead of an automatically allocated buffer. If the *Buffer* parameter is a null character pointer, input/output is completely unbuffered.

The **setbuffer** subroutine is not needed under normal circumstances because the default file I/O buffer size is optimal.

The **setlinebuf** subroutine is used to change the **stdout** or **stderr** file from block buffered or unbuffered to line-buffered. Unlike the **setbuf** and **setbuffer** subroutines, the **setlinebuf** subroutine can be used any time *Stream* is active.

A buffer is normally obtained from the **malloc** subroutine at the time of the first **getc** subroutine or **putc** subroutine on the file, except that the standard error stream, **stderr**, is normally not buffered.

Output streams directed to terminals are always either line-buffered or unbuffered.

**Note:** 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.

The **setbuffer** and **setlinebuf** subroutines are included for compatibility with Berkeley System Distribution (BSD).

## Parameters

<i>Stream</i>	Specifies the input/output stream.
<i>Buffer</i>	Points to a character array.
<i>Mode</i>	Determines how the <i>Stream</i> parameter is buffered.
<i>Size</i>	Specifies the size of the buffer to be used.

## Return Values

Upon successful completion, **setvbuf** returns a value of 0. Otherwise it returns a nonzero value if a value that is not valid is given for type, or if the request cannot be honored.

## Related Information

The **fopen**, **freopen**, or **fdopen** subroutine, **fread** subroutine, **getc**, **fgetc**, **getchar**, or **getw** subroutine, **getwc**, **fgetwc**, or **getwchar** subroutine, **malloc**, **free**, **realloc**, **calloc**, **mallopt**, **mallinfo**, or **alloca** subroutine, **putc**, **putchar**, **fputc**, or **putw** subroutine, **putwc**, **putwchar**, or **fputwc** subroutine.

The Input and Output Handling in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## setcsmap Subroutine

### Purpose

Reads a code-set map file and assigns it to the standard input device.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int setcsmap (Path);  
char * Path;
```

### Description

The **setcsmap** subroutine reads in a code-set map file. The *path* parameter specifies the location of the code-set map file. The path is usually composed by forming a string with the **csmap** directory and the code set, as in the following example:

```
n=sprintf(path,"%s%s",CSMAP_DIR,nl_langinfo(CODESET));
```

The file is processed and according to the included informations, the **setcsmap** subroutine changes the tty configuration. Multibyte processing may be enabled, and converter modules may be pushed onto the tty stream.

## Parameter

*Path* Names the code-set map file.

## Return Values

If a code set-map file is successfully opened and compiled, a value of 0 is returned. If an error occurred, a value of 1 is returned and the **errno** global variable is set to identify the error.

## Error Codes

<b>EINVAL</b>	Indicates an invalid value in the code set map.
<b>EIO</b>	An I/O error occurred while the file system was being read.
<b>ENOMEM</b>	Insufficient resources are available to satisfy the request.
<b>EFAULT</b>	A kernel service, such as <b>copyin</b> , has failed.
<b>ENOENT</b>	The named file does not exist.
<b>EACCESS</b>	The named file cannot be read.

## Related Information

The **setmaps** command.

The **setmaps** file format.

tty Subsystem Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## setgid, setrgid, setegid, setregid, or setgidx Subroutine

### Purpose

Sets the process group IDs.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>

int setgid (GID)
gid_t GID;

int setrgid (RGID)
gid_t RGID;

int setegid (EGID)
gid_t EGID;

int setregid (RGID, EGID)
gid_t RGID;
gid_t EGID;

#include <unistd.h>
#include <sys/id.h>

int setgidx ( which, GID )
int which;
gid_t GID;
```

## Description

The **setgid**, **setrgid**, **setregid**, **setregid**, and **setgidx** subroutines set the process group IDs of the calling process. The following semantics are supported:

<b>setgid</b>	If the effective user ID of the process is the root user, the process's real, effective, and saved group IDs are set to the value of the <i>GID</i> parameter. Otherwise, the process effective group ID is reset if the <i>GID</i> parameter is equal to either the current real or saved group IDs, or one of its supplementary group IDs. Supplementary group IDs of the calling process are not changed.
<b>setegid</b>	The process effective group ID is reset if one of the following conditions is met: <ul style="list-style-type: none"><li>• The <i>EGID</i> parameter is equal to either the current real or saved group IDs.</li><li>• The <i>EGID</i> parameter is equal to one of its supplementary group IDs.</li><li>• The effective user ID of the process is the root user.</li></ul>
<b>setrgid</b>	The <b>EPERM</b> error code is always returned.
<b>setregid</b>	The <i>RGID</i> and <i>EGID</i> parameters can have one of the following relationships: <i>RGID</i> != <i>EGID</i> If the <i>EGID</i> parameter is equal to either the process's real or saved group IDs, the process effective group ID is set to the <i>EGID</i> parameter. Otherwise, the <b>EPERM</b> error code is returned. <i>RGID</i> == <i>EGID</i> If the effective user ID of the process is the root user, the process's real and effective group IDs are set to the <i>EGID</i> parameter. If the <i>EGID</i> parameter is equal to the process's real or saved group IDs, the process effective group ID is set to <i>EGID</i> . Otherwise, the <b>EPERM</b> error code is returned.
<b>setgidx</b>	The <i>which</i> parameter can have one of the following values: <b>ID_EFFECTIVE</b> <i>GID</i> must be either the real or saved <i>GID</i> or one of the values in the concurrent group set. The effective group ID for the current process will be set to <i>GID</i> . <b>ID_EFFECTIVEID_REAL</b> Invoker must have appropriate privilege. The real and effective group ID for the current process will be set to <i>GID</i> . <b>ID_EFFECTIVEID_REALID_SAVED</b> Invoker must have appropriate privilege. The real, effective and saved group ID for the current process will be set to <i>GID</i> .

The **setegid**, **setrgid**, **setregid**, and **setgidx** subroutines are thread-safe.

The operating system does not support **setuid** ("setuid, setruid, seteuid, setreuid or setuidx Subroutine" on page 151) or **setgid** shell scripts.

These subroutines are part of Base Operating System (BOS) Runtime.

## Parameters

<i>GID</i>	Specifies the value of the group ID to set.
<i>RGID</i>	Specifies the value of the real group ID to set.
<i>EGID</i>	Specifies the value of the effective group ID to set.
<i>which</i>	Specifies which group ID values to set.

## Return Values

0	Indicates that the subroutine was successful.
-1	Indicates the subroutine failed. The <b>errno</b> global variable is set to indicate the error.

## Error Codes

If the **setgid**, **setegid**, or **setgidx** subroutine fails, one or more of the following are returned:

- EPERM** Indicates the process does not have appropriate privileges and the *GID* or *EGID* parameter is not equal to either the real or saved group IDs of the process.
- EINVAL** Indicates the value of the *GID*, *EGID* or *which* parameter is invalid.

## Related Information

The **getgid** subroutine, **getgroups** subroutine, **setgroups** (“setgroups Subroutine”) subroutine, **setuid** (“setuid, setruid, seteuid, setreuid or setuidx Subroutine” on page 151) subroutine.

The **setgroups** command.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## setgroups Subroutine

### Purpose

Sets the supplementary group ID of the current process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <grp.h>
```

```
int setgroups ( NumberGroups, GroupIDSet)
int NumberGroups;
gid_t *GroupIDSet;
```

### Description

The **setgroups** subroutine sets the supplementary group ID of the process. The **setgroups** subroutine cannot set more than **NGROUPS\_MAX** groups in the group set. (**NGROUPS\_MAX** is a constant defined in the **limits.h** file.)

**Note:** The routine may coredump instead of returning EFAULT when an invalid pointer is passed in case of 64-bit application calling 32-bit kernel interface.

### Parameters

*GroupIDSet* Pointer to the array of group IDs to be established.

*NumberGroups* Indicates the number of entries in the *GroupIDSet* parameter.

### Return Values

Upon successful completion, the **setgroups** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **setgroups** subroutine fails if any of the following are true:

- EFAULT** The *NumberGroups* and *GroupIDSet* parameters specify an array that is partially or completely outside of the process' allocated address space.
- EINVAL** The *NumberGroups* parameter is greater than the **NGROUPS\_MAX** value.
- EPERM** A group ID in the *GroupIDSet* parameter is not presently in the supplementary group ID, and the invoker does not have root user authority.

## Security

Auditing Events:

Event	Information
<b>PROC_SetGroups</b>	<i>NumberGroups, GroupIDSet</i>

## Related Information

The **getgid** subroutine, **getgroups** subroutine, **initgroups** subroutine, **setgid** ("setgid, setrgid, setegid, setregid, or setgidx Subroutine" on page 132) subroutine.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## setjmp or longjmp Subroutine

### Purpose

Saves and restores the current execution context.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <setjmp.h>
int setjmp (Context)
jmp_buf Context;
```

```
void longjmp ( Context, Value)
jmp_buf Context;
int Value;
```

```
int _setjmp (Context)
jmp_buf Context;
```

```
void _longjmp (Context, Value)
jmp_buf Context;
int Value;
```

### Description

The **setjmp** subroutine and the **longjmp** subroutine are useful when handling errors and interrupts encountered in low-level subroutines of a program.

The **setjmp** subroutine saves the current stack context and signal mask in the buffer specified by the *Context* parameter.

The **longjmp** subroutine restores the stack context and signal mask that were saved by the **setjmp** subroutine in the corresponding *Context* buffer. After the **longjmp** subroutine runs, program execution continues as if the corresponding call to the **setjmp** subroutine had just returned the value of the *Value* parameter. The subroutine that called the **setjmp** subroutine must not have returned before the completion of the **longjmp** subroutine. The **setjmp** and **longjmp** subroutines save and restore the signal mask **sigmask (2)**, while **\_setjmp** and **\_longjmp** manipulate only the stack context.

If a process is using the AT&T System V **sigset** interface, then the **setjmp** and **longjmp** subroutines do not save and restore the signal mask. In such a case, their actions are identical to those of the **\_setjmp** and **\_longjmp** subroutines.

## Parameters

*Context*        Specifies an address for a **jmp\_buf** structure.  
*Value*         Indicates any integer value.

## Return Values

The **setjmp** subroutine returns a value of 0, unless the return is from a call to the **longjmp** function, in which case **setjmp** returns a nonzero value.

The **longjmp** subroutine cannot return 0 to the previous context. The value 0 is reserved to indicate the actual return from the **setjmp** subroutine when first called by the program. The **longjmp** subroutine does not return from where it was called, but rather, program execution continues as if the corresponding call to **setjmp** was returned with a returned value of *Value*.

If the **longjmp** subroutine is passed a *Value* parameter of 0, then execution continues as if the corresponding call to the **setjmp** subroutine had returned a value of 1. All accessible data have values as of the time the **longjmp** subroutine is called.

**Attention:** If the **longjmp** subroutine is called with a *Context* parameter that was not previously set by the **setjmp** subroutine, or if the subroutine that made the corresponding call to the **setjmp** subroutine has already returned, then the results of the **longjmp** subroutine are undefined. If the **longjmp** subroutine detects such a condition, it calls the **longjmperror** routine. If **longjmperror** returns, the program is aborted. The default version of **longjmperror** prints the message: `longjmp or siglongjmp used outside of saved context to standard error` and returns. Users wishing to exit in another manner can write their own version of the **longjmperror** program.

## Related Information

The **sigsetjmp** or **siglongjmp** (“sigsetjmp or siglongjmp Subroutine” on page 186) subroutine.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## setlocale Subroutine

### Purpose

Changes or queries the program’s entire current locale or portions thereof.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <locale.h>
```

```
char *setlocale ( Category, Locale)  
int Category;  
const char *Locale;
```

## Description

The **setlocale** subroutine selects all or part of the program's locale specified by the *Category* and *Locale* parameters. The **setlocale** subroutine then changes or queries the specified portion of the locale. The **LC\_ALL** value for the *Category* parameter names the entire locale (all the categories). The other *Category* values name only a portion of the program locale.

The *Locale* parameter specifies a string that provides information needed to set certain conventions in the *Category* parameter. The components of the *Locale* parameter are language and territory. Values allowed for the locale argument are the predefined **language\_territory** combinations or a user-defined locale.

If a user defines a new locale, a uniquely named locale definition source file must be provided. The character collation, character classification, monetary, numeric, time, and message information should be provided in this file. The locale definition source file is converted to a binary file by the **localedef** command. The binary locale definition file is accessed in the directory specified by the **LOCPATH** environment variable.

**Note:** All **setuid** and **setgid** programs will ignore the **LOCPATH** environment variable.

The default locale at program startup is the C locale. A call to the **setlocale** subroutine must be made explicitly to change this default locale environment. See *Understanding Locale Subroutines in AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs* for **setlocale** subroutine examples.

The locale state is common to all threads within a process.

## Parameters

*Category* Specifies a value representing all or part of the locale for a program. Depending on the value of the *Locale* parameter, these categories may be initiated by the values of environment variables with corresponding names. Valid values for the *Category* parameter, as defined in the **locale.h** file, are:

### **LC\_ALL**

Affects the behavior of a program's entire locale.

### **LC\_COLLATE**

Affects the behavior of regular expression and collation subroutines.

### **LC\_CTYPE**

Affects the behavior of regular expression, character-classification, case-conversion, and wide character subroutines.

### **LC\_MESSAGES**

Affects the content of messages and affirmative and negative responses.

### **LC\_MONETARY**

Affects the behavior of subroutines that format monetary values.

### **LC\_NUMERIC**

Affects the behavior of subroutines that format nonmonetary numeric values.

### **LC\_TIME**

Affects the behavior of time-conversion subroutines.

*Locale* Points to a character string containing the required setting for the *Category* parameter.

The following are special values for the *Locale* parameter:

**"C"** The C locale is the locale all programs inherit at program startup.

**"POSIX"**  
Specifies the same locale as a value of **"C"**.

**""** Specifies categories be set according to locale environment variables.

**NULL** Queries the current locale environment and returns the name of the locale.

The Language Territory Table contains supported **language\_territory** values for the *Locale* parameter:

Table 1. Language Territory Table

<i>Locale Value</i>	<b>Language</b>	<b>Territory</b>	<b>Code Set</b>
<b>Ar_AA</b>	Arabic	Arabic Countries	IBM-1046
<b>ar_AA</b>	Arabic	Arabic Countries	ISO8859-6
<b>bg_BG</b>	Bulgarian	Bulgaria	ISO8856-5
<b>cs_CZ</b>	Czech	Czech Republic	ISO8859-2
<b>da_DK</b>	Danish	Denmark	ISO8859-1
<b>de_CH</b>	German	Switzerland	ISO8859-1
<b>de_DE</b>	German	Germany	ISO8859-1
<b>el_GR</b>	Greek	Greece	ISO8859-7
<b>en_GB</b>	English	Great Britain	ISO8859-1
<b>en_US</b>	English	United States	ISO8859-1
<b>es_ES</b>	Spanish	Spain	ISO8859-1
<b>fi_FI</b>	Finnish	Finland	ISO8859-1
<b>fr_BE</b>	French	Belgium	ISO8859-1
<b>fr_CA</b>	French	Canada	ISO8859-1
<b>fr_FR</b>	French	France	ISO8859-1
<b>fr_CH</b>	French	Switzerland	ISO8859-1
<b>hr_HR</b>	Croatian	Croatia	ISO8859-2
<b>hu_HU</b>	Hungarian	Hungary	ISO8859-2
<b>is_IS</b>	Icelandic	Iceland	ISO8859-1
<b>it_IT</b>	Italian	Italy	ISO8859-1
<b>Iw_IL</b>	Hebrew	Israel	IBM-856
<b>iw_IL</b>	Hebrew	Israel	ISO8859-8
<b>Ja_JP</b>	Japanese	Japan	IBM-943
<b>ja_JP</b>	Japanese	Japan	IBM-eucJP
<b>ko_KR</b>	Korean	Korea	IBM-eucKR
<b>mk_MK</b>	Macedonian	Former Yugoslav Republic of Macedonia	ISO8859-5
<b>nl_BE</b>	Dutch	Belgium	ISO8859-1
<b>nl_NL</b>	Dutch	Netherlands	ISO8859-1
<b>no_NO</b>	Norwegian	Norway	ISO8859-1

Table 1. Language Territory Table (continued)

Locale Value	Language	Territory	Code Set
pl_PL	Polish	Poland	ISO8859-2
pt_PT	Portuguese	Portugal	ISO8859-1
ro_RO	Romanian	Romania	ISO8859-2
ru_RU	Russian	Russia	ISO8859-5
sh_SP	Serbian Latin	Yugoslavia	ISO8859-2
sl_SI	Slovene	Slovenia	ISO8859-2
sk_SK	Slovak	Slovakia	ISO8859-2
sr_SP	Serbian Cyrillic	Yugoslavia	ISO8859-5
Zh_CN	Simplified Chinese	PRC	GBK
sv_SE	Swedish	Sweden	ISO8859-1
tr_TR	Turkish	Turkey	ISO8859-9
zh_TW	Chinese (trad)	Taiwan	IBM-eucTW

## Return Values

If a pointer to a string is given for the *Locale* parameter and the selection can be honored, the **setlocale** subroutine returns the string associated with the specified *Category* parameter for the new locale. If the selection cannot be honored, a null pointer is returned and the program locale is unchanged.

If a null is used for the *Locale* parameter, the **setlocale** subroutine returns the string associated with the *Category* parameter for the program's current locale. The program's locale is not changed.

A subsequent call with the string returned by the **setlocale** subroutine, and its associated category, will restore that part of the program locale. The string returned is not modified by the program, but can be overwritten by a subsequent call to the **setlocale** subroutine.

## Related Information

The **localeconv** subroutine, **nl\_langinfo** subroutine, **rpmatch** ("rpmatch Subroutine" on page 49) subroutine.

The **localedef** command.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Setting the Locale in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## setpcred Subroutine

### Purpose

Sets the current process credentials.

### Library

Security Library (**libc.a**)

## Syntax

```
#include <usersec.h>
```

```
int setpcred ( User, Credentials)  
char **Credentials;  
char *User;
```

## Description

The **setpcred** subroutine sets a process' credentials according to the *Credentials* parameter. If the *User* parameter is specified, the credentials defined for the user in the user database are used. If the *Credentials* parameter is specified, the credentials in this string are used. If both the *User* and *Credentials* parameters are specified, both the user's and the supplied credentials are used. However, the supplied credentials of the *Credentials* parameter will override those of the user. At least one parameter must be specified.

The **setpcred** subroutine requires the **setpenv** subroutine to follow it.

**Note:** If the **auditwrite** subroutine is to be called from a program invoked from the **inittab** file, the **setpcred** subroutine should be called first to establish the process' credentials.

*User*                      Specifies the user for whom credentials are being established.

## Credentials

Defines specific credentials to be established. This parameter points to an array of null-terminated character strings that may contain the following values. The last character string must be null.

**LOGIN\_USER=%s**  
Login user name

**REAL\_USER=%s**  
Real user name

**REAL\_GROUP=%s**  
Real group name

**GROUPS=%s**  
Supplementary group ID

**AUDIT\_CLASSES=%s**  
Audit classes

**RLIMIT\_CPU=%d**  
Process soft CPU limit

**RLIMIT\_FSIZE=%d**  
Process soft file size

**RLIMIT\_DATA=%d**  
Process soft data segment size

**RLIMIT\_STACK=%d**  
Process soft stack segment size

**RLIMIT\_CORE=%d**  
Process soft core file size

**RLIMIT\_RSS=%d**  
Process soft resident set size

**RLIMIT\_CORE\_HARD=%d**  
Process hard core file size

**RLIMIT\_CPU\_HARD=%d**  
Process hard CPU limit

**RLIMIT\_DATA\_HARD=%d**  
Process hard data segment size

**RLIMIT\_FSIZE\_HARD=%d**  
Process hard file size

**RLIMIT\_RSS\_HARD=%d**  
Process hard resident set size

**RLIMIT\_STACK\_HARD=%d**  
Process hard stack segment size

**UMASK=%o**  
Process **umask** (file creation mask)

A process must have root user authority to set all credentials except the UMASK credential.

Resource	Hard	Soft
RLIMIT_CORE	unlimited	%d
RLIMIT_CPU	%d	%d
RLIMIT_DATA	unlimited	%d
RLIMIT_FSIZE	%d	%d
RLIMIT_RSS	unlimited	%d
RLIMIT_STACK	unlimited	%d

The soft limit credentials will override the equivalent hard limit credentials that may proceed them. To set the hard limits, the hard limit credentials should follow the soft limit credentials.

## Return Values

Upon successful return, the **setpcred** subroutine returns a value of 0. If **setpcred** fails, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **setpcred** subroutine fails if one or more of the following are true:

<b>EINVAL</b>	The <i>Credentials</i> parameter contains invalid credentials specifications.
<b>EINVAL</b>	The <i>User</i> parameter is null and the <i>Credentials</i> parameter is either null or points to an empty string.
<b>EPERM</b>	The process does not have the proper authority to set the requested credentials.

Other errors may be set by subroutines invoked by the **setpcred** subroutine.

## Related Information

The **auditwrite** subroutine, **ckuseracct** subroutine, **ckuserID** subroutine, **getpcred** subroutine, **getpenv** subroutine, **setpenv** (“setpenv Subroutine”) subroutine.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## setpenv Subroutine

### Purpose

Sets the current process environment.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
int setpenv ( User, Mode, Environment, Command)
```

```
char *User;
```

```
int Mode;
```

```
char **Environment;
```

```
char *Command;
```

### Description

The **setpenv** subroutine first sets the environment of the current process according to its parameter values, and then sets the working directory and runs a specified command. If the *User* parameter is specified, the process environment is set to that of the specified user, the user's working directory is set, and the specified command run. If the *User* parameter is not specified, then the environment and working directory are set to that of the current process, and the command is run from this process. The environment consists of both user-state and system-state environment variables.

**Note:** The **setpenv** subroutine requires the **setpcred** subroutine to precede it.

The **setpenv** subroutine performs the following steps:

#### Setting the Process Environment

The first step involves changing the user-state and system-state environment. Since this is dependent on the values of the *Mode* and *Environment* parameters, see the description for the *Mode* parameter for more information.

#### Setting the Process Current Working Directory

After the user-state and system-state environment is set, the working directory of the process may be set. If the *Mode* parameter includes the **PENV\_INIT** value, the current working directory is changed to the user's initial login directory (defined in the */etc/passwd* file). Otherwise, the current working directory is unchanged.

#### Executing the Initial Program

After the working directory of the process is reset, the initial program (usually the shell interpreter) is executed. If the *Command* parameter is null, the shell from the user database is used. If the parameter is not defined, the shell from the user-state environment is used and the *Command* parameter defaults to the */usr/bin/sh* file. If the *Command* parameter is not null, it specifies the command to be executed. If the *Mode* parameter contains the **PENV\_ARGV** value, the *Command* parameter is assumed to be in the **argv** structure and is passed to the **execve** subroutine. The string contained in the *Command* parameter is used as the *Path* parameter of the **execve** subroutine. If the *Mode* parameter does not contain **PENV\_ARGV** value, the *Command* parameter is parsed into an **argv** structure and executed. If the *Command* parameter contains the **\$SHELL** value, substitution is done prior to execution.

**Note:** This step will fail if the *Command* parameter contains the **\$SHELL** value but the user-state environment does not contain the **SHELL** value.

## Parameters

#### *Command*

Specifies the command to be executed. If the *Mode* parameter contains the **PENV\_ARGV** value, then the *Command* parameter is assumed to be a valid argument vector for the **execv** subroutine.

#### *Environment*

Specifies the value of user-state and system-state environment variables in the same format returned by the **getpenv** subroutine. The user-state variables are prefaced by the keyword **USRENVIRON:**, and the system-state variables are prefaced by the keyword **SYSENVIRON:**. Each variable is defined by a string of the form **var=value**, which is an array of null-terminated character pointers.

*Mode* Specifies how the **setpenv** subroutine is to set the environment and run the command. This parameter is a bit mask and must contain only one of the following values, which are defined in the **usersec.h** file:

#### **PENV\_INIT**

The user-state environment is initialized as follows:

#### **AUTHSTATE**

Retained from the current environment. If the **AUTHSTATE** value is not present, it is defaulted to the **compat** value.

**KRB5CCNAME**

Retained from the current environment. This value is defined if you authenticated through the Distributed Computing Environment (DCE).

**USER** Set to the name specified by the *User* parameter or to the name corresponding to the current real user ID. The name is shortened to a maximum of 8 characters.

**LOGIN**

Set to the name specified by the *User* parameter or to the name corresponding to the current real user ID. If set by the *User* parameter, this value is the complete login name, which may include a DCE cell name.

**LOGNAME**

Set to the current system environment variable **LOGNAME**.

**TERM** Retained from the current environment. If the **TERM** value is not present, it is defaulted to an **IBM6155**.

**SHELL**

Set from the initial program defined for the real user ID of the current process. If no program is defined, then the **/usr/bin/sh** shell is used as the default.

**HOME** Set from the home directory defined for the real user ID of the current process. If no home directory is defined, the default is **/home/guest**.

**PATH** Set initially to the value for the **PATH** value in the **/etc/environment** file. If not set, it is destructively replaced by the default value of **PATH=/usr/bin:\$HOME:.** (The final period specifies the working directory). The **PATH** variable is destructively replaced by the **usrenv** attribute for this user in the **/etc/security/envIRON** file if the **PATH** value exists in the **/etc/environment** file.

The following files are read for additional environment variables:

**/etc/environment**

Variables defined in this file are added to the environment.

**/etc/security/envIRON**

Environment variables defined for the user in this file are added to the user-state environment.

The user-state variables in the *Environment* parameter are added to the user-state environment. These are preceded by the **USRENVIRON:** keyword.

The system-state environment is initialized as follows:

**LOGNAME**

Set to the current **LOGNAME** value in the protected user environment. The **login** (**tsm**) command passes this value to the **setpenv** subroutine to ensure correctness.

**NAME** Set to the login name corresponding to the real user ID.

**TTY** Set to the TTY name corresponding to standard input.

The following file is read for additional environment variables:

**/etc/security/envIRON**

The system-state environment variables defined for the user in this file are added to the environment. The system-state variables in the *Environment* parameter are added to the environment. These are preceded by the **SYSENVIRON** keyword.

**PENV\_DELTA**

The existing user-state and system-state environment variables are preserved and the variables defined in the *Environment* parameter are added.

**PENV\_RESET**

The existing environment is cleared and totally replaced by the content of the *Environment* parameter.

**PENV\_KLEEN**

Closes all open file descriptors, except 0, 1, and 2, before executing the command. This value must be logically ORed with **PENV\_DELTA**, **PENV\_RESET**, or **PENV\_INIT**. It cannot be used alone.

**PENV\_NOPROF**

The new shell will not be treated as a login shell. Only valid when used with the **PENV\_INIT** flag.

For both system-state and user-state environments, variable substitution is performed.

The *Mode* parameter may also contain:

**PENV\_ARGV** Specifies that the *Command* parameter is already in **argv** format and need not be parsed. This value must be logically ORed with **PENV\_DELTA**, **PENV\_RESET**, or **PENV\_INIT**. It cannot be used alone.

*User* Specifies the user name whose environment and working directory is to be set and the specified command run. If a null pointer is given, the current real uid is used to determine the name of the user.

## Return Values

If the environment was successfully established, this function does not return. If the **setpenv** subroutine fails, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **setpenv** subroutine fails if one or more of the following are true:

**EINVAL** The *Mode* parameter contains values other than **PENV\_INIT**, **PENV\_DELTA**, **PENV\_RESET**, or **PENV\_ARGV**.

**EINVAL** The *Mode* parameter contains more than one of **PENV\_INIT**, **PENV\_DELTA**, or **PENV\_RESET** values.

**EINVAL** The *Environment* parameter is neither null nor empty, and does not contain a valid environment string.

**EPERM** The caller does not have read access to the environment defined for the system, or the user does not have permission to change the specified attributes.

Other errors may be set by subroutines invoked by the **setpenv** subroutine.

## Related Information

The **execl**, **execv**, **execle**, **execve**, **execlp**, **execvp**, or **exec** subroutine, **getpenv** subroutine, **setpcred** (“setpcred Subroutine” on page 139) subroutine.

The **login** command, **su** command.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## setpgid or setpgrp Subroutine

### Purpose

Sets the process group ID.

### Libraries

**setpgid**: Standard C Library (**libc.a**)

**setpgrp**: Standard C Library (**libc.a**);

Berkeley Compatibility Library (**libbsd.a**)

### Syntax

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

```
int setpgid ( ProcessID, ProcessGroupID)
```

```
pid_t ProcessID, ProcessGroupID;
```

```
int setpgrp ( )
```

### Description

The **setpgid** subroutine is used either to join an existing process group or to create a new process group within the session of the calling process. The process group ID of a session leader does not change. Upon return, the process group ID of the process having a process ID that matches the *ProcessID* value is set to the *ProcessGroupID* value. As a special case, if the *ProcessID* value is 0, the process ID of the calling process is used. If *ProcessGroupID* value is 0, the process ID of the indicated process is used.

This function is implemented to support job control.

The **setpgrp** subroutine in the **libc.a** library supports a subset of the function of the **setpgid** subroutine. It has no parameters. It sets the process group ID of the calling process to be the same as its process ID and returns the new value.

In BSD systems, the **setpgrp** subroutine is defined with two parameters, as follows:

```
int setpgrp (ProcessID, ProcessGroup)
```

```
int ProcessID, ProcessGroup;
```

### Parameters

*ProcessID* Specifies the process whose process group ID is to be changed.

*ProcessGroupID* Specifies the new value of calling process group ID.

### Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **setpgid** subroutine is unsuccessful if one or more of the following is true:

- EACCES** The value of the *ProcessID* parameter matches the process ID of a child process of the calling process and the child process has successfully executed one of the **exec** subroutines.
- EINVAL** The value of the *ProcessGroupID* parameter is less than 0, or is not a valid value.
- ENOSYS** The **setpgid** subroutine is not supported by this implementation.
- EPERM** The process indicated by the value of the *ProcessID* parameter is a session leader.
- EPERM** The value of the *ProcessID* parameter matches the process ID of a child process of the calling process and the child process is not in the same session as the calling process.
- EPERM** The value of the *ProcessGroupID* parameter is valid, but does not match the process ID of the process indicated by the *ProcessID* parameter. There is no process with a process group ID that matches the value of the *ProcessGroupID* parameter in the same session as the calling process.
- ESRCH** The value of the *ProcessID* parameter does not match the process ID of the calling process of a child process of the calling process.

## Related Information

The **getpid** subroutine.

---

## setpri Subroutine

### Purpose

Sets a process scheduling priority to a constant value.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int setpri ( ProcessID, Priority)
pid_t ProcessID;
int Priority;
```

### Description

The **setpri** subroutine sets the scheduling priority of all threads in a process to be a constant. All threads have their scheduling policies changed to **SCHED\_RR**. A process nice value and CPU usage can no longer be used to determine a process scheduling priority. Only processes that have root user authority can set a process scheduling priority to a constant.

### Parameters

- ProcessID* Specifies the process ID. If this value is 0 then the current process scheduling priority is set to a constant.
- Priority* Specifies the scheduling priority for the process. A lower number value designates a higher scheduling priority. The *Priority* parameter must be in the range **PRIORITY\_MIN** < *Priority* < **PRIORITY\_MAX**. (See the **sys/sched.h** file.)

### Return Values

Upon successful completion, the **setpri** subroutine returns the former scheduling priority of the process just changed. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **setpri** subroutine is unsuccessful if one or more of the following is true:

- EINVAL** The priority specified by the *Priority* parameter is outside the range of acceptable priorities.
- EPERM** The process executing the **setpri** subroutine call does not have root user authority.
- ESRCH** No process can be found corresponding to that specified by the *ProcessID* parameter.

## Related Information

The **getpri** subroutine.

Performance-Related Subroutines in *AIX 5L Version 5.2 Performance Management Guide*.

---

## setpwdb or endpwdb Subroutine

### Purpose

Opens or closes the authentication database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <userpw.h>

int setpwdb ( Mode)
int Mode;
int endpwdb ( )
```

### Description

These functions are used to open and close access to the authentication database. Programs that call either the **getuserpw** or **putuserpw** subroutine should call the **setpwdb** subroutine to open the database and the **endpwdb** subroutine to close the database.

The **setpwdb** subroutine opens the authentication database in the specified mode, if it is not already open. The open count is increased by 1.

The **endpwdb** subroutine decreases the open count by one and closes the authentication database when this count drops to 0. Subsequent references to individual data items can cause a memory access violation. The **endpwdb** subroutine also frees the space that was allocated by either the **getuserpw**, **putuserpw**, or **putuserpwhist** subroutine. For security reasons, freeing the space clears the password field. Any uncommitted changed data is lost.

### Parameters

*Mode* Specifies the mode of the open. This parameter may contain one or more of the following values, defined in the **usersec.h** file:

**S\_READ**

Specifies read access.

**S\_WRITE**

Specifies update access.

## Return Values

The **setpwdb** and **endpwdb** subroutines return a value of 0 to indicate success. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **setpwdb** and **endpwdb** subroutines fail if the following is true:

**EACCES** Access permission is denied for the data request.

Both of these functions return errors from other subroutines.

## Security

Access Control: The calling process must have access to the authentication data.

Files Accessed:

Modes	File
rw	/etc/security/passwd
rw	/etc/passwd

## Related Information

The **getgroupattr** subroutine, **getuserattr** subroutine, **getuserpw**, **putuserpw**, or **putuserpwhist** subroutine.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## setroledb or endroledb Subroutine

### Purpose

Opens and closes the role database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
int setroledb(Mode)
int Mode;
int endroledb
```

### Description

These functions may be used to open and close access to the role database. Programs that call the **getroleattr** subroutine should call the **setroledb** subroutine to open the role database and the **endroledb** subroutine to close the role database.

The **setroledb** subroutine opens the role database in the specified mode, if it is not already open. The open count is increased by 1.

The **endroledb** subroutine decreases the open count by 1 and closes the role database when this count goes to 0. Any uncommitted changed data is lost.

## Parameters

*Mode* Specifies the mode of the open. This parameter may contain one or more of the following values defined in the **usersec.h** file:

**S\_READ**

Specifies read access.

**S\_WRITE**

Specifies update access.

## Return Values

The **setroledb** and **endroledb** subroutines return a value of 0 to indicate success. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **setroledb** subroutine fails if the following is true:

**EACCES** Access permission is denied for the data request.

Both subroutines return errors from other subroutines.

## Security

Files Accessed: The calling process must have access to the role data.

Mode File **rw/etc/security/roles**

## Related Information

The **getroleattr**, **nextrole**, or **putroleattr** subroutine.

---

## setsid Subroutine

### Purpose

Creates a session and sets the process group ID.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
pid_t setsid (void)
```

### Description

The **setsid** subroutine creates a new session if the calling process is not a process group leader. Upon return, the calling process is the session leader of this new session, the process group leader of a new process group, and has no controlling terminal. The process group ID of the calling process is set equal to its process ID. The calling process is the only process in the new process group and the only process in the new session.

## Return Values

Upon successful completion, the value of the new process group ID is returned. Otherwise, (`pid_t`) -1 is returned and the `errno` global variable is set to indicate the error.

## Error Codes

The `setuid` subroutine is unsuccessful if the following is true:

**EPERM** The calling process is already a process group leader, or the process group ID of a process other than the calling process matches the process ID of the calling process.

## Related Information

The `fork` subroutine, `getpid`, `getpgrp`, or `getppid` subroutine, `setpgid` (“`setpgid` or `setpgrp` Subroutine” on page 146) subroutine, `setpgrp` (“`setpgid` or `setpgrp` Subroutine” on page 146) subroutine.

---

## setuid, setruid, seteuid, setreuid or setuidx Subroutine

### Purpose

Sets the process user IDs.

### Library

Standard C Library (`libc.a`)

### Syntax

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

```
int setuid (UID)
uid_t UID;
```

```
int setruid (RUID)
uid_t RUID;
```

```
int seteuid (EUID)
uid_t EUID;
```

```
int setreuid (RUID, EUID)
uid_t RUID;
uid_t EUID;
```

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

```
int setuidx (which, UID)
int which;
uid_t UID;
```

### Description

The `setuid`, `setruid`, `seteuid`, and `setreuid` subroutines reset the process user IDs. The following semantics are supported:

<b>setuid</b>	If the effective user ID of the process is the root user, the process's real, effective, and saved user IDs are set to the value of the <code>UID</code> parameter. Otherwise, the process effective user ID is reset if the <code>UID</code> parameter specifies either the current real or saved user IDs.
<b>seteuid</b>	The process effective user ID is reset if the <code>UID</code> parameter is equal to either the current real or saved user IDs or if the effective user ID of the process is the root user.
<b>setruid</b>	The <b>EPERM</b> error code is always returned. Processes cannot reset only their real user IDs.

**setreuid** The *RUID* and *EUID* parameters can have the following two possibilities:

*RUID != EUID*  
If the *EUID* parameter specifies either the process's real or saved user IDs, the process effective user ID is set to the *EUID* parameter. Otherwise, the **EPERM** error code is returned.

*RUID == EUID*  
If the process effective user ID is the root user, the process's real and effective user IDs are set to the *EUID* parameter. Otherwise, the **EPERM** error code is returned.

**setuidx** The which parameter can have one of the following values:

**ID\_EFFECTIVE**  
*UID* must be either the real or saved *UID*. The effective user ID for the current process will be set to *UID*.

**ID\_EFFECTIVEID\_REAL**  
Invoker must have appropriate privilege. The real and effective user ID for the current process will be set to *UID*.

**ID\_EFFECTIVEID\_REALID\_SAVED**  
Invoker must have appropriate privilege. The real, effective and saved user ID for the current process will be set to *UID*.

**ID\_LOGIN**  
Invoker must have appropriate privilege. The login UID for the current process will be set to *UID*.

The real and effective user ID parameters can have a value of -1. If the value is -1, the actual value for the *UID* parameter is set to the corresponding current the *UID* parameter of the process.

The operating system does not support **setuid** or **setgid** ("setgid, setrgid, setegid, setregid, or setgidx Subroutine" on page 132) shell scripts.

These subroutines are part of Base Operating System (BOS) Runtime.

## Parameters

*UID* Specifies the user ID to set.  
*EUID* Specifies the effective user ID to set.  
*RUID* Specifies the real user ID to set.  
*which* Specifies which user ID values to set.

## Return Values

Upon successful completion, the **setuid**, **seteuid**, **setreuid**, and **setuidx** subroutines return a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **setuid**, **seteuid**, **setreuid**, and **setuidx** subroutines are unsuccessful if either of the following is true:

**EINVAL** The value of the *UID* or *EUID* parameter is not valid.  
**EPERM** The process does not have the appropriate privileges and the *UID* and *EUID* parameters are not equal to either the real or saved user IDs of the process.

## Related Information

The **getuid** or **geteuid** subroutine, **setgid** ("setgid, setrgid, setegid, setregid, or setgidx Subroutine" on page 132) subroutine.

## setuserdb or enduserdb Subroutine

### Purpose

Opens and closes the user database.

### Library

Security Library (**libc.a**)

### Syntax

```
#include <usersec.h>
```

```
int setuserdb ( Mode)
```

```
int Mode;
```

```
int enduserdb ( )
```

### Description

These functions may be used to open and close access to the user database. Programs that call either the **getuserattr** or **getgroupattr** subroutine should call the **setuserdb** subroutine to open the user database and the **enduserdb** subroutine to close the user database.

The **setuserdb** subroutine opens the user database in the specified mode, if it is not already open. The open count is increased by 1.

The **enduserdb** subroutine decreases the open count by 1 and closes the user database when this count goes to 0. Any uncommitted changed data is lost.

### Parameters

*Mode* Specifies the mode of the open. This parameter may contain one or more of the following values defined in the **usersec.h** file:

**S\_READ**

Specifies read access

**S\_WRITE**

Specifies update access.

### Return Values

The **setuserdb** and **enduserdb** subroutines return a value of 0 to indicate success. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **setuserdb** subroutine fails if the following is true:

**EACCES** Access permission is denied for the data request.

Both subroutines return errors from other subroutines.

## Security

Files Accessed: The calling process must have access to the user data. Depending on the actual attributes accessed, this may include:

Modes	File
rw	/etc/passwd
rw	/etc/group
rw	/etc/security/user
rw	/etc/security/limits
rw	/etc/security/group
rw	/etc/security/environ

## Related Information

The **getgroupattr** subroutine, **getuserattr** subroutine, **getuserpw** subroutine, **setpwdb** (“setpwdb or endpwdb Subroutine” on page 148) subroutine.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## sgetl or sputl Subroutine

### Purpose

Accesses long numeric data in a machine-independent fashion.

### Library

Object File Access Routine Library (**libld.a**)

### Syntax

```
long sgetl ( Buffer )
char *Buffer;
```

```
void sputl ( Value, Buffer )
long Value;
char *Buffer;
```

### Description

The **sgetl** subroutine retrieves four bytes from memory starting at the location pointed to by the *Buffer* parameter. It then returns the bytes as a long *Value* with the byte ordering of the host machine.

The **sputl** subroutine stores the four bytes of the *Value* parameter into memory starting at the location pointed to by the *Buffer* parameter. The order of the bytes is the same across all machines.

Using the **sputl** and **sgetl** subroutines together provides a machine-independent way of storing long numeric data in an ASCII file. For example, the numeric data stored in the portable archive file format can be accessed with the **sputl** and **sgetl** subroutines.

### Parameters

*Value* Specifies a 4-byte value to store into memory.  
*Buffer* Points to a location in memory.

## Related Information

The **ar** command, **dump** command.

The **ar** file format, **a.out** file format.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## shmat Subroutine

### Purpose

Attaches a shared memory segment or a mapped file to the current process.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
void *shmat (SharedMemoryID, SharedMemoryAddress, SharedMemoryFlag)  
int SharedMemoryID, SharedMemoryFlag;  
const void * SharedMemoryAddress;
```

### Description

The **shmat** subroutine attaches the shared memory segment or mapped file specified by the *SharedMemoryID* parameter (returned by the **shmget** subroutine), or file descriptor specified by the *SharedMemoryID* parameter (returned by the **openx** subroutine) to the address space of the calling process.

To learn more about the limits that apply to shared memory, see the Inter-Process Communication (IPC) Limits article in *AIX 5L Version 5.2 General Programming Concepts*.

**Note:** The following applies to AIX 4.2.1 and later releases for 32-bit processes only.

An extended **shmat** capability is available. If an environment variable **EXTSHM=ON** is defined then processes executing in that environment will be able to create and attach more than eleven shared memory segments.

The segments can be of size from 1 byte to 2GB, although for segments larger than 256MB in size the environment variable **EXTSHM=ON** is ignored. The process can attach these segments into the address space for the size of the segment. Another segment could be attached at the end of the first one in the same 256MB segment region. The address at which a process can attach is at page boundaries - a multiple of **SHMLBA\_EXTSHM** bytes. For segments larger than 256MB in size, the address at which a process can attach is at 256MB boundaries, which is a multiple of **SHMLBA** bytes.

The segments can be of size from 1 byte to 256MB. The process can attach these segments into the address space for the size of the segment. Another segment could be attached at the end of the first one in the same 256MB segment region. The address at which a process can attach will be at page boundaries - a multiple of **SHMLBA\_EXTSHM** bytes.

The maximum address space available for shared memory with or without the environment variable and for memory mapping is 2.75GB. An additional segment register "0xE" is available so that the address

space is from 0x30000000 to 0xE0000000. However, a 256MB region starting from 0xD0000000 will be used by the shared libraries and is therefore unavailable for shared memory regions or *mmap*ed regions.

On AIX 5.2 and later, a 32-bit process running with the very large address space model has up to 3.25 GB of address space available for the **shmat** and **mmap** memory mappings. For a 32-bit process with the very large address space model, the address space available for mappings is from 0x30000000 to 0xFFFFFFFF. This extended address range applies to both extended **shmat** and standard **shmat**. For more information on how to use the very large address space model, see the Understanding the Very Large Address-Space Model article in *AIX 5L Version 5.2 General Programming Concepts*.

There are some restrictions on the use of the extended shmat feature. These shared memory regions can not be used as I/O buffers where the unpinning of the buffer occurs in an interrupt handler. The restrictions on the use are the same as that of *mmap* buffers.

The smaller region sizes are not supported for mapping files. Regardless of whether **EXTSHM=ON** or not, mapping a file will consume at least 256MB of address space.

The **SHM\_SIZE shmctl** command is not supported for segments created with **EXTSHM=ON**.

A segment created with **EXTSHM=ON** can be attached by a process without **EXTSHM=ON**. This will consume a 256MB area of the address space irrespective of the size of the shared memory region.

A segment created without **EXTSHM=ON** can be attached by a process with **EXTSHM=ON**. This will consume a 256MB area of the address space irrespective of the size of the shared memory region.

The environment variable provides the option of executing an application either with the additional functionality of attaching more than 11 segments when **EXTSHM=ON**, or the higher-performance access to 11 or fewer segments when the environment variable is not set.

## Parameters

- |                            |  |
|----------------------------|--|
| <i>SharedMemoryID</i>      | Specifies an identifier for the shared memory segment.   |
| <i>SharedMemoryAddress</i> | Identifies the segment or file attached at the address specified by the <i>SharedMemoryAddress</i> parameter, as follows: <ul style="list-style-type: none"><li>• If the <i>SharedMemoryAddress</i> parameter is not equal to 0, and the <b>SHM_RND</b> flag is set in the <i>SharedMemoryFlag</i> parameter, the segment or file is attached at the next lower segment boundary. This address is given by <math>(SharedMemoryAddress - (SharedMemoryAddress \text{ modulo } SHMLBA\_EXTSHM \text{ if environment variable } EXTSHM=ON \text{ or } SHMLBA \text{ if not}))</math>. SHMLBA specifies the low boundary address multiple of a segment.</li><li>• If the <i>SharedMemoryAddress</i> parameter is not equal to 0 and the <b>SHM_RND</b> flag is not set in the <i>SharedMemoryFlag</i> parameter, the segment or file is attached at the address given by the <i>SharedMemoryAddress</i> parameter. If this address does not point to a <b>SHMLBA_EXTSHM</b> boundary if the environment variable <b>EXTSHM=ON</b> or <b>SHMLBA</b> boundary if not, the <b>shmat</b> subroutine returns the value -1 and sets the <b>errno</b> global variable to the <b>EINVAL</b> error code. SHMLBA specifies the low boundary address multiple of a segment.</li></ul> |

<i>SharedMemoryFlag</i>	Specifies several options. Its value is either 0 or is constructed by logically ORing one or more of the following values:
<b>SHM_COPY</b>	Changes an open file to deferred update (see the <b>openx</b> subroutine). Included only for compatibility with previous versions of the operating system.
<b>SHM_MAP</b>	Maps a file onto the address space instead of a shared memory segment. The <i>SharedMemoryID</i> parameter must specify an open file descriptor in this case.
<b>SHM_RDONLY</b>	Specifies read-only mode instead of the default read-write mode.
<b>SHM_RND</b>	Rounds the address given by the <i>SharedMemoryAddress</i> parameter to the next lower segment boundary, if necessary.

The **shmat** subroutine makes a shared memory segment addressable by the current process. The segment is attached for reading if the **SHM\_RDONLY** flag is set and the current process has read permission. If the **SHM\_RDONLY** flag is not set and the current process has both read and write permission, it is attached for reading and writing.

If the **SHM\_MAP** flag is set, file mapping takes place. In this case, the **shmat** subroutine maps the file open on the file descriptor specified by the *SharedMemoryID* onto a segment. The file must be a regular file. The segment is then mapped into the address space of the process. A file of any size can be mapped if there is enough space in the user address space.

When file mapping is requested, the *SharedMemoryFlag* parameter specifies how the file should be mapped. If the **SHM\_RDONLY** flag is set, the file is mapped read-only. To map read-write, the file must have been opened for writing.

All processes that map the same file read-only or read-write map to the same segment. This segment remains mapped until the last process mapping the file closes it.

A mapped file opened with the **O\_DEFER** update has deferred update. That is, changes to the shared segment do not affect the contents of the file resident in the file system until an **fsync** subroutine is issued to the file descriptor for which the mapping was requested. Setting the **SHM\_COPY** flag changes the file to the deferred state. The file remains in this state until all processes close it. The **SHM\_COPY** flag is provided only for compatibility with Version 2 of the operating system. New programs should use the **O\_DEFER** open flag.

A file descriptor can be used to map the corresponding file only once. To map a file several times requires multiple file descriptors.

When a file is mapped onto a segment, the file is referenced by accessing the segment. The memory paging system automatically takes care of the physical I/O. References beyond the end of the file cause the file to be extended in page-sized increments. The file cannot be extended beyond the next segment boundary.

## Return Values

When successful, the segment start address of the attached shared memory segment or mapped file is returned. Otherwise, the shared memory segment is not attached, the **errno** global variable is set to indicate the error, and a value of -1 is returned.

## Error Codes

The **shmat** subroutine is unsuccessful and the shared memory segment or mapped file is not attached if one or more of the following are true:

<b>EACCES</b>	The calling process is denied permission for the specified operation.
<b>EAGAIN</b>	The file to be mapped has enforced locking enabled, and the file is currently locked.
<b>EBADF</b>	A file descriptor to map does not refer to an open regular file.
<b>EEXIST</b>	The file to be mapped has already been mapped.
<b>EINVAL</b>	The <b>SHM_RDONLY</b> and <b>SHM_COPY</b> flags are both set.
<b>EINVAL</b>	The <i>SharedMemoryID</i> parameter is not a valid shared memory identifier.
<b>EINVAL</b>	The <i>SharedMemoryAddress</i> parameter is not equal to 0, and the value of ( <i>SharedMemoryAddress</i> - ( <i>SharedMemoryAddress</i> modulo <b>SHMLBA_EXTSHM</b> if the environment variable <b>EXTSHM=ON</b> or <b>SHMLBA</b> if not ) points outside the address space of the process.
<b>EINVAL</b>	The <i>SharedMemoryAddress</i> parameter is not equal to 0, the <b>SHM_RND</b> flag is not set in the <i>SharedMemoryFlag</i> parameter, and the <i>SharedMemoryAddress</i> parameter points to a location outside of the address space of the process.
<b>EMFILE</b>	The number of shared memory segments attached to the calling process exceeds the system-imposed limit.
<b>ENOMEM</b>	The available data space in memory is not large enough to hold the shared memory segment. ENOMEM is always returned if a 32-bit process tries to attach a shared memory segment larger than 2GB.
<b>ENOMEM</b>	The available data space in memory is not large enough to hold the mapped file data structure.
<b>ENOMEM</b>	The requested address and length crosses a segment boundary. This is not supported when the environment variable <b>EXTSHM=ON</b> .

## Related Information

The **exec** subroutine, **exit** subroutine, **fclean** subroutine, **fork** subroutine, **fsync** subroutine, **mmap** subroutine, **munmap** subroutine, **openx** subroutine, **truncate** subroutine, **readvx** subroutine, **shmctl** subroutine, **shmdt** subroutine, **shmget** subroutine, **writevx** subroutine.

The **ipcs** command and **ipcrm** command.

List of Memory Manipulation Services, Subroutines Overview, Understanding Memory Mapping in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## shmctl Subroutine

### Purpose

Controls shared memory operations.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int shmctl (SharedMemoryID, Command, Buffer)
int SharedMemoryID, Command;
struct shmids * Buffer;
```

### Description

The **shmctl** subroutine performs a variety of shared-memory control operations as specified by the *Command* parameter.

The following limits apply to shared memory:

- Maximum shared-memory segment size is:
  - 256M bytes before AIX 4.3.1
  - 2G bytes for AIX 4.3.1 through AIX 5.1
  - 64G bytes for 64-bit applications for AIX 5.1 and later
- Minimum shared-memory segment size is 1 byte.
- Maximum number of shared memory IDs is 4096 for operating system releases before AIX 4.3.2 and 131072 for AIX 4.3.2 and following.

## Parameters

*SharedMemoryID*

Specifies an identifier returned by the **shmget** subroutine.

*Buffer*

Indicates a pointer to the **shmid\_ds** structure. The **shmid\_ds** structure is defined in the **sys/shm.h** file.

*Command*

The following commands are available:

### IPC\_STAT

Obtains status information about the shared memory segment identified by the *SharedMemoryID* parameter. This information is stored in the area pointed to by the *Buffer* parameter. The calling process must have read permission to run this command.

### IPC\_SET

Sets the user and group IDs of the owner as well as the access permissions for the shared memory segment identified by the *SharedMemoryID* parameter. This command sets the following fields:

```
shm_perm.uid /* owning user ID      */
shm_perm.gid /* owning group ID     */
shm_perm.mode /* permission bits only */
```

You must have an effective user ID equal to root or to the value of the *shm\_perm.cuid* or *shm\_perm.uid* field in the **shmid\_ds** data structure identified by the *SharedMemoryID* parameter.

### IPC\_RMID

Removes the shared memory identifier specified by the *SharedMemoryID* parameter from the system and erases the shared memory segment and data structure associated with it. This command is only executed by a process that has an effective user ID equal either to that of superuser or to the value of the *shm\_perm.uid* or *shm\_perm.cuid* field in the data structure identified by the *SharedMemoryID* parameter.

### SHM\_SIZE

Sets the size of the shared memory segment to the value specified by the *shm\_segsz* field of the structure specified by the *Buffer* parameter. This value can be larger or smaller than the current size. The limit is the maximum shared-memory segment size. This command is only executed by a process that has an effective user ID equal either to that of a process with the appropriate privileges or to the value of the *shm\_perm.uid* or *shm\_perm.cuid* field in the data structure identified by the *SharedMemoryID* parameter. This command is not supported for regions created with the environment variable **EXTSHM=ON**. This results in a return value of -1 with **errno** set to **EINVAL**. Attempting to use the **SHM\_SIZE** on a shared memory region larger than 256MB or attempting to increase the size of a shared memory region larger than 256MB results in a return value of -1 with **errno** set to **EINVAL**.

## Return Values

When completed successfully, the **shmctl** subroutine returns a value of 0. Otherwise, it returns a value of -1 and the **errno** global variable is set to indicate the error.

## Error Codes

The **shmctl** subroutine is unsuccessful if one or more of the following are true:

- EACCES** The *Command* parameter is equal to the **IPC\_STAT** value and read permission is denied to the calling process.
- EFAULT** The *Buffer* parameter points to a location outside the allocated address space of the process.
- EINVAL** The *SharedMemoryID* parameter is not a valid shared memory identifier.
- EINVAL** The *Command* parameter is not a valid command.
- EINVAL** The *Command* parameter is equal to the **SHM\_SIZE** value and the value of the *shm\_segsz* field of the structure specified by the *Buffer* parameter is not valid.
- EINVAL** The *Command* parameter is equal to the **SHM\_SIZE** value and the shared memory region was created with the environment variable **EXTSHM=ON**.
- ENOMEM** The *Command* parameter is equal to the **SHM\_SIZE** value, and the attempt to change the segment size is unsuccessful because the system does not have enough memory.
- EOVERFLOW** The *Command* parameter is **IPC\_STAT** and the size of the shared memory region is greater than or equal to 4G bytes. This only happens with 32-bit programs.
- EPERM** The *Command* parameter is equal to the **IPC\_RMID** or **SHM\_SIZE** value, and the effective user ID of the calling process is not equal to the value of the *shm\_perm.uid* or *shm\_perm.cuid* field in the data structure identified by the *SharedMemoryID* parameter. The effective user ID of the calling process is not the root user ID.

## Related Information

The **disclaim** subroutine, **shmat** subroutine, **shmdt** subroutine, **shmget** subroutine.

The **ipcs** command and **ipcrm** command.

List of Memory Manipulation Services, Subroutines Overview, Understanding Memory Mapping in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## shmdt Subroutine

### Purpose

Detaches a shared memory segment.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int shmdt (SharedMemoryAddress)  
const void * SharedMemoryAddress;
```

### Description

The **shmdt** subroutine detaches from the data segment of the calling process the shared memory segment located at the address specified by the *SharedMemoryAddress* parameter.

Mapped file segments are automatically detached when the mapped file is closed. However, you can use the **shmdt** subroutine to explicitly release the segment register used to map a file. Shared memory segments must be explicitly detached with the **shmdt** subroutine.

If the file was mapped for writing, the **shmdt** subroutine updates the **mtime** and **ctime** time stamps.

The following limits apply to shared memory:

- Maximum shared-memory segment size is:
  - 256M bytes before AIX 4.3.1
  - 2G bytes for AIX 4.3.1 through AIX 5.1
  - 64G bytes for 64-bit applications for AIX 5.1 and later
- Minimum shared-memory segment size is 1 byte.
- Maximum number of shared memory IDs is 4096 for operating system releases before AIX 4.3.2 and 131072 for AIX 4.3.2 and following.

## Parameters

*SharedMemoryAddress* Specifies the data segment start address of a shared memory segment.

## Return Values

When successful, the **shmdt** subroutine returns a value of 0. Otherwise, the shared memory segment at the address specified by the *SharedMemoryAddress* parameter is not detached, a value of 1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **shmdt** subroutine is unsuccessful if the following condition is true:

**EINVAL** The value of the *SharedMemoryAddress* parameter is not the data-segment start address of a shared memory segment.

## Related Information

The **exec** subroutine, **exit** subroutine, **fork** subroutine, **fsync** subroutine, **mmap** subroutine, **munmap** subroutine, **shmat** subroutine, **shmctl** subroutine, **shmget** subroutine.

The **ipcs** command and **ipcrm** command.

List of Memory Manipulation Services, Subroutines Overview, Understanding Memory Mapping in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## shmget Subroutine

### Purpose

Gets shared memory segments.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int shmget (Key, Size, SharedMemoryFlag)
```

```
key_t Key;
```

```
size_t Size
```

```
int SharedMemoryFlag;
```

## Description

The **shmget** subroutine returns the shared memory identifier associated with the specified *Key* parameter.

The following limits apply to shared memory:

- Maximum shared-memory segment size is:
  - 256M bytes before AIX 4.3.1
  - 2G bytes for AIX 4.3.1 through AIX 5.1
  - 64G bytes for 64-bit applications for AIX 5.1 and later
- Minimum shared-memory segment size is 1 byte.
- Maximum number of shared memory IDs is 4096 for operating system releases before AIX 4.3.2 and 131072 for AIX 4.3.2 and following.

## Parameters

*Key* Specifies either the **IPC\_PRIVATE** value or an IPC key constructed by the **ftok** subroutine (or by a similar algorithm).

*Size* Specifies the number of bytes of shared memory required.

*SharedMemoryFlag* Constructed by logically ORing one or more of the following values:

### **IPC\_CREAT**

Creates the data structure if it does not already exist.

### **IPC\_EXCL**

Causes the **shmget** subroutine to be unsuccessful if the **IPC\_CREAT** flag is also set, and the data structure already exists.

### **SHM\_LGPAGE**

Attempts to create the region so it can be mapped through hardware-supported, large-page mechanisms, if enabled. This is purely advisory. For the system to consider this flag, it must be used in conjunction with the **SHM\_PIN** flag and enabled with the **vm tune** command (-L to reserve memory for the region (which requires a reboot) and -S to enable **SHM\_PIN**). To successfully get large-pages, the user requesting large-page shared memory must have **CAP\_BYPASS\_RAC\_VMM** capability. This has no effect on shared memory regions created with the **EXTSHM=ON** environment variable.

### **SHM\_PIN**

Attempts to pin the shared memory region if enabled. This is purely advisory. For the system to consider this flag, the system must be enable with **vm tune** command. This has no effect on shared memory regions created with **EXTSHM=ON** environment variable.

### **S\_IRUSR**

Permits the process that owns the data structure to read it.

### **S\_IWUSR**

Permits the process that owns the data structure to modify it.

### **S\_IRGRP**

Permits the group associated with the data structure to read it.

### **S\_IWGRP**

Permits the group associated with the data structure to modify it.

### **S\_IROTH**

Permits others to read the data structure.

### **S\_IWOTH**

Permits others to modify the data structure.

Values that begin with the **S\_I** prefix are defined in the **sys/mode.h** file and are a subset of the access permissions that apply to files.

A shared memory identifier, its associated data structure, and a shared memory segment equal in number of bytes to the value of the *Size* parameter are created for the *Key* parameter if one of the following is true:

- The *Key* parameter is equal to the **IPC\_PRIVATE** value.
- The *Key* parameter does not already have a shared memory identifier associated with it, and the **IPC\_CREAT** flag is set in the *SharedMemoryFlag* parameter.

Upon creation, the data structure associated with the new shared memory identifier is initialized as follows:

- The *shm\_perm.cuid* and *shm\_perm.uid* fields are set to the effective user ID of the calling process.
- The *shm\_perm.cgid* and *shm\_perm.gid* fields are set to the effective group ID of the calling process.
- The low-order 9 bits of the *shm\_perm.mode* field are set to the low-order 9 bits of the *SharedMemoryFlag* parameter.
- The *shm\_segsz* field is set to the value of the *Size* parameter.
- The *shm\_lpid*, *shm\_nattch*, *shm\_atime*, and *shm\_dtime* fields are set to 0.
- The *shm\_ctime* field is set to the current time.

**Note:** Once created, a shared memory segment is deleted only when the system reboots or by issuing the **ipcrm** command or using the following **shmctl** subroutine:

```
if (shmctl (id, IPC_RMID, 0) == -1)
    perror ("error in closing segment"),exit (1);
```

## Return Values

Upon successful completion, a shared memory identifier is returned. Otherwise, the **shmget** subroutine returns a value of -1 and sets the **errno** global variable to indicate the error.

## Error Codes

The **shmget** subroutine is unsuccessful if one or more of the following are true:

- |               |   |
|---------------|---|
| <b>EACCES</b> | A shared memory identifier exists for the <i>Key</i> parameter, but operation permission as specified by the low-order 9 bits of the <i>SharedMemoryFlag</i> parameter is not granted.                  |
| <b>EEXIST</b> | A shared memory identifier exists for the <i>Key</i> parameter, and both the <b>IPC_CREAT</b> and <b>IPC_EXCL</b> flags are set in the <i>SharedMemoryFlag</i> parameter.                               |
| <b>EINVAL</b> | A shared memory identifier does not exist and the <i>Size</i> parameter is less than the system-imposed minimum or greater than the system-imposed maximum.   |
| <b>EINVAL</b> | A shared memory identifier exists for the <i>Key</i> parameter, but the size of the segment associated with it is less than the <i>Size</i> parameter, and the <i>Size</i> parameter is not equal to 0. |
| <b>ENOENT</b> | A shared memory identifier does not exist for the <i>Key</i> parameter, and the <b>IPC_CREAT</b> flag is not set in the <i>SharedMemoryFlag</i> parameter.  |
| <b>ENOMEM</b> | A shared memory identifier and associated shared memory segment are to be created but the amount of available physical memory is not sufficient to meet the request.                                    |
| <b>ENOSPC</b> | A shared memory identifier will be created, but the system-imposed maximum of shared memory identifiers allowed will be exceeded.   |

## Related Information

The **ftok** subroutine, **mmap** subroutine, **munmap** subroutine, **shmat** subroutine, **shmctl** subroutine, **shmdt** subroutine.

The **ipcs** command and **ipcrm** command.

List of Memory Manipulation Services, Subroutines Overview, Understanding Memory Mapping in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## sigaction, sigvec, or signal Subroutine

### Purpose

Specifies the action to take upon delivery of a signal.

### Libraries

**sigaction** Standard C Library (**libc.a**)  
**signal, sigvec** Standard C Library (**libc.a**);

Berkeley Compatibility Library (**libbsd.a**)

### Syntax

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

```
int sigaction ( Signal, Action, OAction)  
int Signal;  
struct sigaction *Action, *OAction;
```

```
int sigvec (Signal, Invec, Outvec)  
int Signal;  
struct sigvec *Invec, *Outvec;
```

```
void (*signal (Signal, Action)) ()  
int Signal;  
void (*Action) (int);
```

### Description

The **sigaction** subroutine allows a calling process to examine and change the action to be taken when a specific signal is delivered to the process issuing this subroutine.

In multi-threaded applications using the threads library (**libpthreads.a**), signal actions are common to all threads within the process. Any thread calling the **sigaction** subroutine changes the action to be taken when a specific signal is delivered to the threads process, that is, to any thread within the process.

**Note:** The **sigaction** subroutine must not be used concurrently to the **sigwait** subroutine on the same signal.

The *Signal* parameter specifies the signal. If the *Action* parameter is not null, it points to a **sigaction** structure that describes the action to be taken on receipt of the *Signal* parameter signal. If the *OAction* parameter is not null, it points to **asigaction** structure in which the signal action data in effect at the time of the **sigaction** subroutine call is returned. If the *Action* parameter is null, signal handling is unchanged; thus, the call can be used to inquire about the current handling of a given signal.

The **sigaction** structure has the following fields:

Member Type	Member Name	Description
void(*) (int)	sa_handler	SIG_DFL, SIG_IGN or pointer to a function.
sigset_t	sa_mask	Additional set of signals to be blocked during execution of signal-catching function.

Member Type	Member Name	Description
int	sa_flags	Special flags to affect behaviour of signal.
void(*) (int, siginfo_t *, void *)	sa_sigaction	Signal-catching function.

The `sa_handler` field can have a **SIG\_DFL** or **SIG\_IGN** value, or it can be a pointer to a function. A **SIG\_DFL** value requests default action to be taken when a signal is delivered. A value of **SIG\_IGN** requests that the signal have no effect on the receiving process. A pointer to a function requests that the signal be caught; that is, the signal should cause the function to be called. These actions are more fully described in "Parameters".

When a signal is delivered to a thread, if the action of that signal specifies termination, stop, or continue, the entire process is terminated, stopped, or continued, respectively.

If the **SA\_SIGINFO** flag (see below) is cleared in the `sa_flags` field of the **sigaction** structure, the `sa_handler` field identifies the action to be associated with the specified signal. If the **SA\_SIGINFO** flag is set in the `sa_flags` field, the `sa_sigaction` field specifies a signal-catching function. If the **SA\_SIGINFO** bit is cleared and the `sa_handler` field specifies a signal-catching function, or if the **SA\_SIGINFO** bit is set, the `sa_mask` field identifies a set of signals that will be added to the signal mask of the thread before the signal-catching function is invoked.

The `sa_mask` field can be used to specify that individual signals, in addition to those in the process signal mask, be blocked from being delivered while the signal handler function specified in the `sa_handler` field is operating. The `sa_flags` field can have the **SA\_ONSTACK**, **SA\_OLDSTYLE**, or **SA\_NOCLDSTOP** bits set to specify further control over the actions taken on delivery of a signal.

If the **SA\_ONSTACK** bit is set, the system runs the signal-catching function on the signal stack specified by the **sigstack** subroutine. If this bit is not set, the function runs on the stack of the process to which the signal is delivered.

If the **SA\_OLDSTYLE** bit is set, the signal action is set to **SIG\_DFL** label prior to calling the signal-catching function. This is supported for compatibility with old applications, and is not recommended since the same signal can recur before the signal-catching subroutine is able to reset the signal action and the default action (normally termination) is taken in that case.

If a signal for which a signal-catching function exists is sent to a process while that process is executing certain subroutines, the call can be restarted if the **SA\_RESTART** bit is set for each signal. The only affected subroutines are the following:

- **read, readx, readv, or readvx** ("read, readx, readv, readvx, or pread Subroutine" on page 16)
- **write, writex, writev, or writevx** ("write, writex, writev, writevx or pwrite Subroutines" on page 506)
- **ioctl orioctlx**
- **fcntl, lockf, or flock**
- **wait, wait3, orwaitpid** ("wait, waitpid, wait3, or wait364 Subroutine" on page 440)

Other subroutines do not restart and return **EINTR** label, independent of the setting of the **SA\_RESTART** bit.

If **SA\_SIGINFO** is cleared and the signal is caught, the signal-catching function will be entered as: `void func(int signo);`

Where *signo* is the only argument to the signal catching function. In this case the **sa\_handler** member must be used to describe the signal catching function and the application must not modify the

**sa\_sigaction** member. If **SA\_SIGINFO** is set and the signal is caught, the signal-catching function will be entered as: `void func(int signo, siginfo_t * info, void * context);` where two additional arguments are passed to the signal catching function.

The second argument will point to an object of type **siginfo\_t** explaining the reason why the signal was generated. The third argument can be cast to a pointer to an object of type **ucontext\_t** to refer to the receiving process' context that was interrupted when the signal was delivered. In this case the **sa\_sigaction** member must be used to describe the signal catching function and the application must not modify the **sa\_handler** member.

The **si\_signo** member contains the system-generated signal number. The **si\_errno** member may contain implementation-dependent additional error information. If nonzero, it contains an error number identifying the condition that caused the signal to be generated. The **si\_code** member contains a code identifying the cause of the signal. If the value of **si\_code** is less than or equal to **0**, the signal was generated by a process and **si\_pid** and **si\_uid** respectively indicate the process ID and the real user ID of the sender.

The **signal.h** header description contains information about the signal specific contents of the elements of the **siginfo\_t** type. If **SA\_NOCLDWAIT** is set and **sig** equals **SIGCHLD**, child processes of the calling processes will not be transformed into zombie processes when they terminate. If the calling process subsequently waits for its children, and the process has no unwaited for children that were transformed into zombie processes, it will block until all of its children terminate, and **wait**, **wait3**, **waitid** and **waitpid** will fail and set **errno** to **ECHILD**. Otherwise, terminating child processes will be transformed into zombie processes, unless **SIGCHLD** is set to **SIG\_IGN**.

If **SA\_RESETHAND** is set, the disposition of the signal will be reset to **SIG\_DFL** and the **SA\_SIGINFO** flag will be cleared on entry to the signal handler.

If **SA\_NODEFER** is set and *sig* is caught, *sig* will not be added to the process' signal mask on entry to the signal handler unless it is included in **sa\_mask**. Otherwise, *sig* will always be added to the process' signal mask on entry to the signal handler. If *sig* is **SIGCHLD**, the **SA\_NOCLDSTOP** flag is not set in **sa\_flags**, and the implementation supports the **SIGCHLD** signal, a **SIGCHLD** signal will be generated for the calling process whenever any of its child processes stop.

If *sig* is **SIGCHLD** and the **SA\_NOCLDSTOP** flag is set in **sa\_flags**, the implementation will not generate a **SIGCHLD** signal in this way. When a signal is caught by a signal-catching function installed by **sigaction**, a new signal mask is calculated and installed for the duration of the signal-catching function (or until a call to either **sigprocmask** or **sigsuspend** is made).

This mask is formed by taking the union of the current signal mask and the value of the **sa\_mask** for the signal being delivered unless **SA\_NODEFER** or **SA\_RESETHAND** is set, and including the signal being delivered. If the user's signal handler returns normally, the original signal mask is restored.

Once an action is installed for a specific signal, it remains installed until another action is explicitly requested (by another call to **sigaction**), until the **SA\_RESETHAND** flag causes resetting of the handler, or until one of the **exec** functions is called.

If the previous action for *sig* had been established by **signal**, the values of the fields returned in the structure pointed to by *oact* are unspecified, and in particular **oact->sa\_handler** is not necessarily the same value passed to **signal**.

However, if a pointer to the same structure or a copy thereof is passed to a subsequent call to **sigaction** through the *act* argument, handling of the signal will be as if the original call to **signal** were repeated.

If **sigaction** fails, no new signal handler is installed. It is unspecified whether an attempt to set the action for a signal that cannot be caught or ignored to **SIG\_DFL** is ignored or causes an error to be returned with **errno** set to **EINVAL**.

If **SA\_SIGINFO** is not set in **sa\_flags**, then the disposition of subsequent occurrences of sig when it is already pending is implementation-dependent; the signal-catching function will be invoked with a single argument.

The **sigvec** and **signal** subroutines are provided for compatibility to older operating systems. Their function is a subset of that available with **sigaction**.

The **sigvec** subroutine uses the **sigvec** structure instead of the **sigaction** structure. The **sigvec** structure specifies a mask as an int instead of a sigset\_t. The mask for the **sigvec** subroutine is constructed by setting the i-th bit in the mask if signal i is to be blocked. Therefore, the **sigvec** subroutine only allows signals between the values of 1 and 31 to be blocked when a signal-handling function is called. The other signals are not blocked by the signal-handler mask.

The **sigvec** structure has the following members:

```
int (*sv_handler)();
/* signal handler */
int sv_mask;
/* signal mask */
int sv_flags;
/* flags */
```

The **sigvec** subroutine in the **libbsd.a** library interprets the **SV\_INTERRUPT** flag and inverts it to the **SA\_RESTART** flag of the **sigaction** subroutine. The **sigvec** subroutine in the **libc.a** library always sets the **SV\_INTERRUPT** flag regardless of what was passed in the **sigvec** structure.

The **sigvec** subroutine in the **libbsd.a** library interprets the **SV\_INTERRUPT** flag and inverts it to the **SA\_RESTART** flag of the **sigaction** subroutine. The **sigvec** subroutine in the **libc.a** library always sets the **SV\_INTERRUPT** flag regardless of what was passed in the **sigvec** structure.

The **signal** subroutine in the **libc.a** library allows an action to be associated with a signal. The *Action* parameter can have the same values that are described for the sv\_handler field in the **sigaction** structure of the **sigaction** subroutine. However, no signal handler mask or flags can be specified; the **signal** subroutine implicitly sets the signal handler mask to additional signals and the flags to be **SA\_OLDSTYLE**.

Upon successful completion of a **signal** call, the value of the previous signal action is returned. If the call fails, a value of -1 is returned and the **errno** global variable is set to indicate the error as in the **sigaction** call.

The **signal** in **libc.a** does not set the **SA\_RESTART** flag. It sets the signal mask to the signal whose action is being specified, and sets flags to **SA\_OLDSTYLE**. The Berkeley Software Distribution (BSD) version of **signal** sets the **SA\_RESTART** flag and preserves the current settings of the signal mask and flags. The BSD version can be used by compiling with the Berkeley Compatibility Library (**libbsd.a**).

The **signal** in **libc.a** does not set the **SA\_RESTART** flag. It sets the signal mask to the signal whose action is being specified, and sets flags to **SA\_OLDSTYLE**. The Berkeley Software Distribution (BSD) version of **signal** sets the **SA\_RESTART** flag and preserves the current settings of the signal mask and flags. The BSD version can be used by compiling with the Berkeley Compatibility Library (**libbsd.a**).

## Parameters

*Signal* Defines the signal. The following list describes signal names and the specification for each. The value of the *Signal* parameter can be any signal name from this list or its corresponding number except the **SIGKILL** name. If you use the signal name, you must include the **signal.h** file, because the name is correlated in the file with its corresponding number.

**Note:** The symbols in the following list of signals represent these actions:

- \* Specifies the default action that includes creating a core dump file.
- @ Specifies the default action that stops the process receiving these signals.
- ! Specifies the default action that restarts or continues the process receiving these signals.
- + Specifies the default action that ignores these signals.
- % Indicates a likely shortage of paging space.
- # See *Terminal Programming* for more information on the use of these signals.

**SIGHUP**

Hang-up. (1)

**SIGINT**

Interrupt. (2)

**SIGQUIT**

Quit. (3\*)

**SIGILL**

Invalid instruction (not reset when caught). (4\*)

**SIGTRAP**

Trace trap (not reset when caught). (5\*)

**SIGIOT**

End process (see the **abort** subroutine). (6\*)

**SIGEMT**

EMT instruction. (7\*)

**SIGFPE**

Arithmetic exception, integer divide by 0, or floating-point exception.(8\*)

**SIGKILL**

Kill (cannot be caught or ignored). (9)

**SIGBUS**

Specification exception. (10\*)

**SIGSEGV**

Segmentation violation. (11\*)

**SIGSYS**

Parameter not valid to subroutine. (12\*)

**SIGPIPE**

Write on a pipe when there is no process to read it. (13)

**SIGALRM**

Alarm clock. (14)

**SIGTERM**

Software termination signal. (15)

**SIGURG**

Urgent condition on I/O channel. (16+)

**SIGSTOP**

Stop (cannot be caught or ignored). (17@)

**SIGTSTP**

Interactive stop. (18@)

**SIGCONT**

Continue if stopped. (19!)

**SIGCHLD**  
To parent on child stop or exit. (20+)

**SIGTTIN**  
Background read attempted from control terminal. (21@)

**SIGTTOU**  
Background write attempted from control terminal. (22@)

**SIGIO** Input/output possible or completed. (23+)

**SIGXCPU**  
CPU time limit exceeded (see the **setrlimit** subroutine). (24)

**SIGXFSZ**  
File size limit exceeded (see the **setrlimit** subroutine).(25)

**reserved**  
(26)

**SIGMSG**  
Input data has been stored into the input ring buffer. (27#)

**SIGWINCH**  
Window size change. (28+)

**SIGPWR**  
Power-fail restart. (29+)

**SIGUSR1**  
User-defined signal 1. (30)

**SIGUSR2**  
User-defined signal 2. (31)

**SIGPROF**  
Profiling timer expired. (see the **setitimer** subroutine).(32)

**SIGDANGER**  
Paging space low. (33+%)

**SIGVTALRM**  
Virtual time alarm (see the **setitimer** subroutine). (34)

**SIGMIGRATE**  
Migrate process. (35)

**SIGPRE**  
Programming exception (user defined). (36)

**reserved**  
(37-58)

**SIGGRANT**  
Monitor access wanted. (60#)

**SIGRETRACT**  
Monitor access should be relinquished. (61#)

**SIGSOUND**  
A sound control has completed execution. (62#)

**SIGSAK**  
Secure attention key. (63)

*Action* Points to a **sigaction** structure that describes the action to be taken upon receipt of the *Signal* parameter signal.

The three types of actions that can be associated with a signal (**SIG\_DFL**, **SIG\_IGN**, or a pointer to a function) are described as follows:

- **SIG\_DFL** Default action: signal-specific default action.

Except for those signal numbers marked with a + (plus sign), @ (at sign), or ! (exclamation point), the default action for a signal ends the receiving process with all of the consequences described in the `_exit` subroutine. In addition, a memory image file is created in the current directory of the receiving process if an asterisk appears with a *Signal* parameter and the following conditions are met:

- The saved user ID and the real user ID of the receiving process are equal.
- An ordinary file named **core** exists in the current directory and is writable, or it can be created. If the file is created, it must have the following properties:
  - The access permission code 0666 (0x1B6), modified by the file-creation mask (see the **umask** subroutine)
  - A file owner ID that is the same as the effective user ID of the receiving process
  - A file group ID that is the same as the effective group ID of the receiving process.

For signal numbers marked with a ! (exclamation point), the default action restarts the receiving process if it has stopped, or continues to run the receiving process.

For signal numbers marked with a @ (at sign), the default action stops the execution of the receiving process temporarily. When a process stops, a **SIGCHLD** signal is sent to its parent process, unless the parent process has set the **SA\_NOCLDSTOP** bit. While a process has stopped, any additional signals that are sent are not delivered until the process has started again. An exception to this is the **SIGKILL** signal, which always terminates the receiving process. Another exception is the **SIGCONT** signal, which always causes the receiving process to restart or continue running. A process whose parent process has ended is sent a **SIGKILL** signal if the **SIGTSTP**, **SIGTTIN**, or **SIGTTOU** signals are generated for that process.

For signal numbers marked with a +, the default action ignores the signal. In this case, the delivery of a signal does not affect the receiving process.

If a signal action is set to **SIG\_DFL** while the signal is pending, the signal remains pending.

- **SIG\_IGN** Ignore signal.

Delivery of the signal does not affect the receiving process. If a signal action is set to the **SIG\_IGN** action while the signal is pending, the pending signal is discarded.

An exception to this is the **SIGCHLD** signal whose **SIG\_DFL** action ignores the signal. If the action for the **SIGCHLD** signal is set to **SIG\_IGN**, child processes of the calling processes will not be transformed into zombie processes when they terminate. If the calling process subsequently waits for its children, and the process has no unwaited for children that were transformed into zombie processes, it will block until all of its children terminate, and **wait**, **wait3**, **waitid** and **waitpid** will fail and set **errno** to **ECHILD**.

**Note:** The **SIGKILL** and **SIGSTOP** signals cannot be ignored.

- Pointer to a function, catch signal.

Upon delivery of the signal, the receiving process runs the signal-catching function specified by the pointer to function. The signal-handler subroutine can be declared as follows:

```
handler(Signal, Code, SCP)
int Signal, Code;
struct sigcontext *SCP;
```

The *Signal* parameter is the signal number. The *Code* parameter is provided only for compatibility with other UNIX-compatible systems. The *Code* parameter value is always 0. The *SCP* parameter points to the **sigcontext** structure that is later used to restore the previous execution context of the process. The **sigcontext** structure is defined in the **signal.h** file.

A new signal mask is calculated and installed for the duration of the signal-catching function (or until **sigprocmask** or **sigsuspend** subroutine is made). This mask is formed by joining the process-signal mask (the mask associated with the action for the signal being delivered) and the mask corresponding to the signal being delivered. The mask associated with the signal-catching function is not allowed to block those signals that cannot be ignored. This is enforced by the kernel without causing an error to be indicated. If and when the signal-catching function returns, the original signal mask is restored (modified by any **sigprocmask** calls that were made since the signal-catching function was called) and the receiving process resumes execution at the point it was interrupted.

The signal-catching function can cause the process to resume in a different context by calling the **longjmp** subroutine. When the **longjmp** subroutine is called, the process leaves the signal stack, if it is currently on the stack, and restores the process signal mask to the state when the corresponding **setjmp** subroutine was made.

Once an action is installed for a specific signal, it remains installed until another action is explicitly requested (by another call to the **sigaction** subroutine), or until one of the **exec** subroutines is called. An exception to this is when the **SA\_OLDSTYLE** bit is set. In this case the action of a caught signal gets set to the **SIG\_DFL** action before the signal-catching function for that signal is called.

If a signal action is set to a pointer to a function while the signal is pending, the signal remains pending.

When signal-catching functions are invoked asynchronously with process execution, the behavior of some of the functions defined by this standard is unspecified if they are called from a signal-catching function. The following set of functions are reentrant with respect to signals; that is, applications can invoke them, without restriction, from signal-catching functions:

**\_exit**

**access**

**alarm**

**cfgetispeed**

**cfgetospeed**

**cfsetispeed**

**cfsetospeed**

**chdir**

**chmod**

**chown**

**close**

**creat**

**dup**

**dup2**

**exec**

**execle**  
**execve**  
**fcntl**  
**fork**  
**fpathconf**  
**fstat**  
**getegid**  
**geteuid**  
**getgid**  
**getgroups**  
**getpgrp**  
**getpid**  
**getppid**  
**getuid**  
**kill**  
**link**  
**lseek**  
**mkdir**  
**mkfifo**  
**open**  
**pathconf**  
**pause**  
**pipe**  
**raise**  
**read**  
**readx**  
**rename**  
**rmdir**  
**setgid**  
**setpgid**  
**setpgrp**  
**setsid**  
**setuid**  
**sigaction**  
**sigaddset**  
**sigdelset**  
**sigemptyset**

**sigismember**  
**signal**  
**sigpending**  
**sigprocmask**  
**sigsuspend**  
**sleep**  
**stat**  
**statx**  
**sysconf**  
**tcdrain**  
**tcflow**  
**tcflush**  
**tcgetattr**  
**tcgetpgrp**  
**tcsendbreak**  
**tcsetattr**  
**tcsetpgrp**  
**time**  
**times**  
**umask**  
**uname**  
**unlink**  
**ustat**  
**utime**  
**wait**  
**waitpid**  
**write**

All other subroutines should not be called from signal-catching functions since their behavior is undefined.

*OAction*

Points to a **sigaction** structure in which the signal action data in effect at the time of the **sigaction** subroutine is returned.

*Invec* Points to a **sigvec** structure that describes the action to be taken upon receipt of the *Signal* parameter signal.

*Outvec*

Points to a **sigvec** structure in which the signal action data in effect at the time of the **sigvec** subroutine is returned.

*Action* Specifies the action associated with a signal.

## Return Values

Upon successful completion, the **sigaction** subroutine returns a value of 0. Otherwise, a value of **SIG\_ERR** is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **sigaction** subroutine is unsuccessful and no new signal handler is installed if one of the following occurs:

- EFAULT** The *Action* or *OAction* parameter points to a location outside of the allocated address space of the process.
- EINVAL** The *Signal* parameter is not a valid signal number.
- EINVAL** An attempt was made to ignore or supply a handler for the **SIGKILL**, **SIGSTOP**, and **SIGCONT** signals.

## Related Information

The **acct** subroutine, **\_exit**, **exit**, or **atexit** subroutine, **getinterval**, **incinterval**, **absinterval**, **resinc**, **resabs**, **alarm**, **ualarm**, **getitimer**, or **setitimer** subroutine, **getrlimit**, **setrlimit**, or **vlimit** subroutine, **kill** subroutine, **longjmp** or **setjmp** (“setjmp or longjmp Subroutine” on page 135) subroutine, **pause** subroutine, **ptrace** subroutine, **sigpause** or **sigsuspend** (“sigsuspend or sigpause Subroutine” on page 188) subroutine, **sigprocmask**, **sigsetmask**, or **sigblock** (“sigprocmask, sigsetmask, or sigblock Subroutine” on page 179) subroutine, **sigstack** (“sigstack Subroutine” on page 187) subroutine, **sigwait** (“sigwait Subroutine” on page 192) subroutine, **umask** (“umask Subroutine” on page 418) subroutine, **wait**, **waitpid**, or **wait3** (“wait, waitpid, wait3, or wait364 Subroutine” on page 440) subroutine.

The **kill** command.

The **core** file.

Signal Management in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs* provides more information about signal management in multi-threaded processes.

---

## sigaltstack Subroutine

### Purpose

Allows a thread to define and examine the state of an alternate stack for signal handlers.

### Library

(**libc.a**)

### Syntax

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

```
int sigaltstack(const stack_t *ss, stack_t *oss);
```

### Description

The **sigaltstack** subroutine allows a thread to define and examine the state of an alternate stack for signal handlers. Signals that have been explicitly declared to execute on the alternate stack will be delivered on the alternate stack.

If *ss* is not null pointer, it points to a **stack\_t** structure that specifies the alternate signal stack that will take effect upon return from **sigaltstack** subroutine. The **ss\_flags** member specifies the new stack state. If it is

set to **SS\_DISABLE**, the stack is disabled and **ss\_sp** and **ss\_ssize** are ignored. Otherwise the stack will be enabled, and the **ss\_sp** and **ss\_size** members specify the new address and size of the stack.

The range of addresses starting at **ss\_sp**, up to but not including **ss\_sp + ss\_size**, is available to the implementation for use as the stack.

If *oss* is not a null pointer, on successful completion it will point to a **stack\_t** structure that specifies the alternate signal stack that was in effect prior to the **sigaltstack** subroutine. The **ss\_sp** and **ss\_size** members specify the address and size of the stack. The **ss\_flags** member specifies the stack's state, and may contain one of the following values:

- SS\_ONSTACK**      The process is currently executing on the alternate signal stack. Attempts to modify the alternate signal stack while the process is executing or it fails. This flag must not be modified by processes.
- SS\_DISABLE**      The alternate signal stack is currently disabled.

The value of **SIGSTKSZ** is a system default specifying the number of bytes that would be used to cover the usual case when manually allocating an alternate stack area. The value **MINSIGSTKSZ** is defined to be the minimum stack size for a signal handler. In computing an alternate stack size, a program should add that amount to its stack requirements to allow for the system implementation overhead.

After a successful call to one of the exec functions, there are no alternate stacks in the new process image.

## Parameters

- ss*      A pointer to a **stack\_t** structure specifying the alternate stack to use during signal handling.
- oss*      A pointer to a **stack\_t** structure that will indicate the alternate stack currently in use.

## Return Values

Upon successful completion, **sigaltstack** subroutine returns 0. Otherwise, it returns -1 and set **errno** to indicate the error.

- 1      Not successful and the **errno** global variable is set to one of the following error codes.

## Error Codes

- EINVAL**      The *ss* parameter is not a null pointer, and the **ss\_flags** member pointed to by *ss* contains flags other than **SS\_DISABLE**.
- ENOMEM**      The size of the alternate stack area is less than **MINSIGSTKSZ**.
- EPERM**      An attempt was made to modify an active stack.

## Related Information

The **sigaction** ("sigaction, sigvec, or signal Subroutine" on page 164) subroutine, **sigsetjmp** ("sigsetjmp or siglongjmp Subroutine" on page 186) subroutine.

---

## sigemptyset, sigfillset, sigaddset, sigdelset, or sigismember Subroutine

### Purpose

Creates and manipulates signal masks.

## Library

Standard C Library (**libc.a**)

## Syntax

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

```
int sigemptyset ( Set)  
sigset_t *Set;
```

```
int sigfillset (Set)  
sigset_t *Set;
```

```
int sigaddset (Set, SignalNumber)  
sigset_t *Set;  
int SignalNumber;
```

```
int sigdelset (Set, SignalNumber)  
sigset_t *Set;  
int SignalNumber;
```

```
int sigismember (Set, SignalNumber)  
sigset_t *Set;  
int SignalNumber;
```

## Description

The **sigemptyset**, **sigfillset**, **sigaddset**, **sigdelset**, and **sigismember** subroutines manipulate sets of signals. These functions operate on data objects addressable by the application, not on any set of signals known to the system, such as the set blocked from delivery to a process or the set pending for a process.

The **sigemptyset** subroutine initializes the signal set pointed to by the *Set* parameter such that all signals are excluded. The **sigfillset** subroutine initializes the signal set pointed to by the *Set* parameter such that all signals are included. A call to either the **sigfillset** or **sigemptyset** subroutine must be made at least once for each object of the **sigset\_t** type prior to any other use of that object.

The **sigaddset** and **sigdelset** subroutines respectively add and delete the individual signal specified by the *SignalNumber* parameter from the signal set specified by the *Set* parameter. The **sigismember** subroutine tests whether the *SignalNumber* parameter is a member of the signal set pointed to by the *Set* parameter.

## Parameters

<i>Set</i>	Specifies the signal set.
<i>SignalNumber</i>	Specifies the individual signal.

## Examples

To generate and use a signal mask that blocks only the **SIGINT** signal from delivery, enter:

```
#include <signal.h>  
  
int return_value;  
sigset_t newset;  
sigset_t *newset_p;  
.  
.  
.  
newset_p = &newset;  
sigemptyset(newset);  
sigaddset(newset, SIGINT);  
return_value = sigprocmask (SIG_SETMASK, newset_p, NULL);
```

## Return Values

Upon successful completion, the **sigismember** subroutine returns a value of 1 if the specified signal is a member of the specified set, or the value of 0 if not. Upon successful completion, the other subroutines return a value of 0. For all the preceding subroutines, if an error is detected, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **sigfillset**, **sigdelset**, **sigismember**, and **sigaddset** subroutines are unsuccessful if the following is true:

**EINVAL**     The value of the *SignalNumber* parameter is not a valid signal number.

## Related Information

The **sigaction**, **sigvec**, or **signal** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine, **sigprocmask** (“sigprocmask, sigsetmask, or sigblock Subroutine” on page 179) subroutine, **sigsuspend** (“sigsuspend or sigpause Subroutine” on page 188) subroutine.

---

## siginterrupt Subroutine

### Purpose

Sets restart behavior with respect to signals and subroutines.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int siginterrupt ( Signal, Flag )
    int Signal, Flag;
```

### Description

The **siginterrupt** subroutine is used to change the subroutine restart behavior when a subroutine is interrupted by the specified signal. If the flag is false (0), subroutines are restarted if they are interrupted by the specified signal and no data has been transferred yet.

If the flag is true (1), the restarting of subroutines is disabled. If a subroutine is interrupted by the specified signal and no data has been transferred, the subroutine will return a value of -1 with the **errno** global variable set to **EINTR**. Interrupted subroutines that have started transferring data return the amount of data actually transferred. Subroutine interrupt is the signal behavior found on 4.1 BSD and AT&T System V UNIX systems.

Note that the BSD signal-handling semantics are not altered in any other way. Most notably, signal handlers always remain installed until explicitly changed by a subsequent **sigaction** or **sigvec** call, and the signal mask operates as documented in the **sigaction** subroutine. Programs can switch between restartable and interruptible subroutine operations as often as desired in the running of a program.

Issuing a **siginterrupt** call during the running of a signal handler causes the new action to take place on the next signal caught.

Restart does not occur unless it is explicitly specified with the **sigaction** or **sigvec** subroutine in the **libc.a** library.

This subroutine uses an extension of the **sigvec** subroutine that is not available in the BSD 4.2; hence, it should not be used if backward compatibility is needed.

## Parameters

*Signal*        Indicates the signal.  
*Flag*         Indicates true or false.

## Return Values

A value of 0 indicates that the call succeeded. A value of -1 indicates that the supplied signal number is not valid.

## Related Information

The **sigaction** or **sigvec** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine, **sigpause** (“sigsuspend or sigpause Subroutine” on page 188) subroutine, **sigsetmask** or **sigblock** (“sigprocmask, sigsetmask, or sigblock Subroutine” on page 179) subroutine.

---

## signbit Macro

### Purpose

Tests the sign.

### Syntax

```
#include <math.h>
```

```
int signbit (x)  
real-floating x;
```

### Description

The **signbit** macro determines whether the sign of its argument value is negative. NaNs, zeros, and infinities have a sign bit.

### Parameters

*x*                Specifies the value to be tested.

### Return Values

The **signbit** macro returns a nonzero value if the sign of its argument value is negative.

### Related Information

class, \_class, finite, isnan, or unordered Subroutines, fpclassify Subroutine, isfinite Subroutine, isinf Subroutine, isnormal Subroutine, and lldiv Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

**math.h** in *AIX 5L Version 5.2 Files Reference*.

---

## sigpending Subroutine

### Purpose

Returns a set of signals that are blocked from delivery.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int sigpending ( Set)  
sigset_t *Set;
```

### Description

The **sigpending** subroutine stores a set of signals that are blocked from delivery and pending for the calling thread, in the space pointed to by the *Set* parameter.

### Parameters

*Set* Specifies the set of signals.

### Return Values

Upon successful completion, the **sigpending** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **sigpending** subroutine is unsuccessful if the following is true:

**EINVAL** The input parameter is outside the user's address space.

### Related Information

The **sigprocmask** ("sigprocmask, sigsetmask, or sigblock Subroutine") subroutine.

---

## sigprocmask, sigsetmask, or sigblock Subroutine

### Purpose

Sets the current signal mask.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int sigprocmask ( How, Set, OSet)
int How;
const sigset_t *Set;
sigset *OSet;
```

```
int sigsetmask ( SignalMask)
int SignalMask;
```

```
int sigblock (SignalMask)
int SignalMask;
```

## Description

**Note:** The **sigprocmask**, **sigsetmask**, and **sigblock** subroutines must not be used in a multi-threaded application. The **sigthreadmask** (“sigthreadmask Subroutine” on page 189) subroutine must be used instead.

The **sigprocmask** subroutine is used to examine or change the signal mask of the calling thread.

The subroutine is used to examine or change the signal mask of the calling process.

Typically, you should use the **sigprocmask(SIG\_BLOCK)** subroutine to block signals during a critical section of code. Then use the **sigprocmask(SIG\_SETMASK)** subroutine to restore the mask to the previous value returned by the **sigprocmask(SIG\_BLOCK)** subroutine.

If there are any pending unblocked signals after the call to the **sigprocmask** subroutine, at least one of those signals will be delivered before the **sigprocmask** subroutine returns.

The **sigprocmask** subroutine does not allow the **SIGKILL** or **SIGSTOP** signal to be blocked. If a program attempts to block either signal, the **sigprocmask** subroutine gives no indication of the error.

## Parameters

<i>How</i>	Indicates the manner in which the set is changed. It can have one of the following values: <b>SIG_BLOCK</b> The resulting set is the union of the current set and the signal set pointed to by the <i>Set</i> parameter. <b>SIG_UNBLOCK</b> The resulting set is the intersection of the current set and the complement of the signal set pointed to by the <i>Set</i> parameter. <b>SIG_SETMASK</b> The resulting set is the signal set pointed to by the <i>Set</i> parameter.
<i>Set</i>	Specifies the signal set. If the value of the <i>Set</i> parameter is not null, it points to a set of signals to be used to change the currently blocked set. If the value of the <i>Set</i> parameter is null, the value of the <i>How</i> parameter is not significant and the process signal mask is unchanged. Thus, the call can be used to inquire about currently blocked signals.
<i>OSet</i>	If the <i>OSet</i> parameter is not the null value, the signal mask in effect at the time of the call is stored in the space pointed to by the <i>OSet</i> parameter.
<i>SignalMask</i>	Specifies the signal mask of the process.

## Compatibility Interfaces

The **sigsetmask** subroutine allows changing the process signal mask for signal values 1 to 31. This same function can be accomplished for all values with the **sigprocmask(SIG\_SETMASK)** subroutine. The signal of value *i* will be blocked if the *i*th bit of *SignalMask* parameter is set.

Upon successful completion, the **sigsetmask** subroutine returns the value of the previous signal mask. If the subroutine fails, a value of -1 is returned and the **errno** global variable is set to indicate the error as in the **sigprocmask** subroutine.

The **sigblock** subroutine allows signals with values 1 to 31 to be logically ORed into the current process signal mask. This same function can be accomplished for all values with the **sigprocmask(SIG\_BLOCK)** subroutine. The signal of value *i* will be blocked, in addition to those currently blocked, if the *i*-th bit of the *SignalMask* parameter is set.

It is not possible to block a **SIGKILL** or **SIGSTOP** signal using the **sigblock** or **sigsetmask** subroutine. This restriction is *silently* imposed by the system without causing an error to be indicated.

Upon successful completion, the **sigblock** subroutine returns the value of the previous signal mask. If the subroutine fails, a value of -1 is returned and the **errno** global variable is set to indicate the error as in the **sigprocmask** subroutine.

## Return Values

Upon completion, a value of 0 is returned. If the **sigprocmask** subroutine fails, the signal mask of the process is unchanged, a value of -1 is returned, and the global variable **errno** is set to indicate the error.

## Error Codes

The **sigprocmask** subroutine is unsuccessful if the following is true:

<b>EPERM</b>	The user does not have the privilege to change the signal's mask.
<b>EINVAL</b>	The value of the <i>How</i> parameter is not equal to one of the defined values.
<b>EFAULT</b>	The user's mask is not in the process address space.

## Examples

To set the signal mask to block only the **SIGINT** signal from delivery, enter:

```
#include <signal.h>

int return_value;
sigset_t newset;
sigset_t *newset_p;
. . .
newset_p = &newset;
sigemptyset(newset_p);
sigaddset(newset_p, SIGINT);
return_value = sigprocmask (SIG_SETMASK, newset_p, NULL);
```

## Related Information

The **kill** or **killpg** subroutine, **sigaction**, **sigvec**, or **signal** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine, **sigaddset**, **sigdelset**, **sigemptyset**, **sigfillset**, **sigismember** (“sigemptyset, sigfillset, sigaddset, sigdelset, or sigismember Subroutine” on page 175) subroutine, **sigpause** (“sigsuspend or sigpause Subroutine” on page 188) subroutine, **sigpending** (“sigpending Subroutine” on page 179) subroutine, **sigsuspend** (“sigsuspend or sigpause Subroutine” on page 188) subroutine.

---

## sigqueue Subroutine

### Purpose

Queues a signal to a process.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <signal.h>

int sigqueue (pid, signo, value)
pid_t pid;
int signo;
const union sigval value;
```

## Description

The **sigqueue** subroutine causes the signal specified by the *signo* parameter to be sent with the value specified by the *value* parameter to the process specified by the *pid* parameter. If the *signo* parameter is zero, error checking is performed but no signal is actually sent. This can be used to check the validity of the *pid* parameter.

The conditions required for a process to have permission to queue a signal to another process are the same as for the **kill** subroutine.

The **sigqueue** subroutine returns immediately. If **SA\_SIGINFO** is set by the receiving process for the specified signal, and if the resources are available to queue the signal, the signal is queued and sent to the receiving process. If **SA\_SIGINFO** is not set for the *signo* parameter, the signal is sent at least once to the receiving process.

If multiple signals in the range **SIGRTMIN** to **SIGRTMAX** should be available for delivery, the lowest numbered of them will be delivered first.

## Parameters

*pid* Specifies the process to which a signal is to be sent.  
*signo* Specifies the signal number.  
*value* Specifies the value to be sent with the signal.

## Return Values

Upon successful completion the **sigqueue** subroutine returns a zero. If unsuccessful, it returns a -1 and sets the **errno** variable to indicate the error.

## Error Code

The **sigqueue** subroutine will fail if:

<b>EAGAIN</b>	No resources are available to queue the signal. The process has already queued <b>SIGQUEUE_MAX</b> signals that are still pending at the receiver(s), or a system-wide resource limit has been exceeded.
<b>EINVAL</b>	The value of the <i>signo</i> parameter is an invalid or unsupported signal number, or if the selected signal can either stop or continue the receiving process. AIX does not support queuing of the following signals: SIGKILL, SIGSTOP, SIGTSTP, SIGCONT, SIGTTIN, SIGTTOU, and SIGCLD.
<b>EPERM</b>	The process does not have the appropriate privilege to send the signal to the receiving process.
<b>ESRCH</b>	The process specified by the <i>pid</i> parameter does not exist.

## Related Information

“sigtimedwait and sigwaitinfo Subroutine” on page 191 and “sigaction, sigvec, or signal Subroutine” on page 164.

---

## sigset, sighold, sigrelse, or sigignore Subroutine

### Purpose

Enhance the signal facility and provide signal management.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <signal.h>
void (*sigset( Signal, Function))()
int Signal;
void (*Function)();
int sighold ( Signal)
int Signal;
int sigrelse ( Signal)
int Signal;
int sigignore ( Signal)
int Signal;
```

### Description

The **sigset**, **sighold**, **sigrelse**, and **sigignore** subroutines enhance the signal facility and provide signal management for application processes.

The **sigset** subroutine specifies the system signal action to be taken upon receiving a *Signal* parameter.

The **sighld** and **sigrelse** subroutines establish critical regions of code. A call to the **sighold** subroutine is analogous to raising the priority level and deferring or holding a signal until the priority is lowered by **sigrelse**. The **sigrelse** subroutine restores the system signal action to the action that was previously specified by the **sigset** structure.

The **sigignore** subroutine sets the action for the *Signal* parameter to **SIG\_IGN**.

The other signal management routine, **signal**, should not be used in conjunction with these routines for a particular signal type.

## Parameters

*Signal* Specifies the signal. The *Signal* parameter can be assigned any one of the following signals:

**SIGHUP**

Hang up

**SIGINT** Interrupt

**SIGQUIT**

Quit\*

**SIGILL** Illegal instruction (not reset when caught)\*

**SIGTRAP**

Trace trap (not reset when caught)\*

**SIGABRT**

Abort\*

**SIGFPE**

Floating point exception\*, or arithmetic exception, integer divide by 0

**SIGSYS**

Bad argument to routine\*

**SIGPIPE**

Write on a pipe with no one to read it

**SIGALRM**

Alarm clock

**SIGTERM**

Software termination signal

**SIGUSR1**

User-defined signal 1

**SIGUSR2**

User-defined signal 2.

\* The default action for these signals is an abnormal termination.

For portability, application programs should use or catch only the signals listed above. Other signals are hardware-dependant and implementation-dependant and may have very different meanings or results across systems. For example, the System V signals (**SIGEMT**, **SIGBUS**, **SIGSEGV**, and **SIGIOT**) are implementation-dependent and are not listed above. Specific implementations may have other implementation-dependent signals.

*Function* Specifies the choice. The *Function* parameter is declared as a type pointer to a function returning void. The *Function* parameter is assigned one of four values: **SIG\_DFL**, **SIG\_IGN**, **SIG\_HOLD**, or an *address* of a signal-catching function. Definitions of the actions taken by each of the values are:

#### **SIG\_DFL**

Terminate process upon receipt of a signal.

Upon receipt of the signal specified by the *Signal* parameter, the receiving process is to be terminated with all of the consequences outlined in the **\_exit** subroutine. In addition, if *Signal* is one of the signals marked with an asterisk above, implementation-dependent abnormal process termination routines, such as a core dump, can be invoked.

#### **SIG\_IGN**

Ignore signal.

Any pending signal specified by the *Signal* parameter is discarded. A pending signal is a signal that has occurred but for which no action has been taken. The system signal action is set to ignore future occurrences of this signal type.

#### **SIG\_HOLD**

Hold signal.

The signal specified by the *Signal* parameter is to be held. Any pending signal of this type remains held. Only one signal of each type is held.

*address* Catch signal.

Upon receipt of the signal specified by the *Signal* parameter, the receiving process is to execute the signal-catching function pointed to by the *Function* parameter. Any pending signal of this type is released. This address is retained across calls to the other signal management functions, **sighold** and **sigrelse**. The signal number *Signal* is passed as the only argument to the signal-catching function. Before entering the signal-catching function, the value of the *Function* parameter for the caught signal is set to **SIG\_HOLD**. During normal return from the signal-catching handler, the system signal action is restored to the *Function* parameter and any held signal of this type is released. If a nonlocal goto (see the **setjmp** subroutine) is taken, the **sigrelse** subroutine must be invoked to restore the system signal action and to release any held signal of this type.

Upon return from the signal-catching function, the receiving process will resume execution at the point at which it was interrupted, except for implementation-defined signals in which this may not be true.

When a signal to be caught occurs during a nonatomic operation such as a call to the **read**, **write**, **open**, or **ioctl** subroutine on a slow device (such as a terminal); during a **pause** subroutine; during a **wait** subroutine that does not return immediately, the signal-catching function is executed. The interrupted routine then returns a value of -1 to the calling process with the **errno** global variable set to **EINTR**.

## Return Values

Upon successful completion, the **sigset** subroutine returns the previous value of the system signal action for the specified *Signal*. Otherwise, it returns **SIG\_ERR** and the **errno** global variable is set to indicate the error.

For the **sighold**, **sigrelse**, and **sigignore** subroutines, a value of 0 is returned upon success. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **sigset**, **sighold**, **sigrelse**, or **sigignore** subroutine is unsuccessful if the following is true:

**EINVAL** The *Signal* value is either an illegal signal number, or the default handling of *Signal* cannot be changed.

## Related Information

The **exit** subroutine, **kill** subroutine, **setjmp** (“setjmp or longjmp Subroutine” on page 135) subroutine, **signal** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine, **wait** (“wait, waitpid, wait3, or wait364 Subroutine” on page 440) subroutine.

---

## sigsetjmp or siglongjmp Subroutine

### Purpose

Saves or restores stack context and signal mask.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <setjmp.h>
```

```
int sigsetjmp ( Environment, SaveMask)
sigjmp_buf Environment;
int SaveMask;
```

```
void siglongjmp (Environment, Value)
sigjmp_buf Environment;
int Value;
```

### Description

The **sigsetjmp** subroutine saves the current stack context, and if the value of the *SaveMask* parameter is not 0, the **sigsetjmp** subroutine also saves the current signal mask of the process as part of the calling environment.

The **siglongjmp** subroutine restores the saved signal mask only if the *Environment* parameter was initialized by a call to the **sigsetjmp** subroutine with a nonzero *SaveMask* parameter argument.

### Parameters

<i>Environment</i>	Specifies an address for a <b>sigjmp_buf</b> structure.
<i>SaveMask</i>	Specifies the flag used to determine if the signal mask is to be saved.
<i>Value</i>	Specifies the return value from the <b>siglongjmp</b> subroutine.

### Return Values

The **sigsetjmp** subroutine returns a value of 0. The **siglongjmp** subroutine returns a nonzero value.

## Related Information

The **setjmp** or **longjmp** (“setjmp or longjmp Subroutine” on page 135) subroutine, **sigaction** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine, **sigprocmask** (“sigprocmask, sigsetmask, or sigblock Subroutine” on page 179) subroutine, **sigsuspend** (“sigsuspend or sigpause Subroutine” on page 188) subroutine.

---

## sigstack Subroutine

### Purpose

Sets and gets signal stack context.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int sigstack ( InStack, OutStack)  
struct sigstack *InStack, *OutStack;
```

### Description

The **sigstack** subroutine defines an alternate stack on which signals are to be processed.

When a signal occurs and its handler is to run on the signal stack, the system checks to see if the process is already running on that stack. If so, it continues to do so even after the handler returns. If not, the signal handler runs on the signal stack, and the original stack is restored when the handler returns.

Use the **sigvec** or **sigaction** subroutine to specify whether a given signal-handler routine is to run on the signal stack.

**Attention:** A signal stack does not automatically increase in size as a normal stack does. If the stack overflows, unpredictable results can occur.

### Parameters

*InStack* Specifies the stack pointer of the new signal stack.

If the value of the *InStack* parameter is nonzero, it points to a **sigstack** structure, which has the following members:

```
caddr_t ss_sp;  
int ss_onstack;
```

The value of *InStack*->*ss\_sp* specifies the stack pointer of the new signal stack. Since stacks grow from numerically greater addresses to lower ones, the stack pointer passed to the **sigstack** subroutine should point to the numerically high end of the stack area to be used.

*InStack*->*ss\_onstack* should be set to a value of 1 if the process is currently running on that stack; otherwise, it should be a value of 0.

*OutStack* If the value of the *InStack* parameter is 0 (that is, a null pointer), the signal stack state is not set. Points to structure where current signal stack state is stored.

If the value of the *OutStack* parameter is nonzero, it points to a **sigstack** structure into which the **sigstack** subroutine stores the current signal stack state.

If the value of the *OutStack* parameter is 0, the previous signal stack state is not reported.

### Return Values

Upon successful completion, the **sigstack** subroutine returns a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **sigstack** subroutine is unsuccessful and the signal stack context remains unchanged if the following is true:

**EFAULT** The *InStack* or *OutStack* parameter points outside of the address space of the process.

## Related Information

The **longjmp** (“setjmp or longjmp Subroutine” on page 135) subroutine, **setjmp** (“setjmp or longjmp Subroutine” on page 135) subroutine, **sigaction**, **signal**, or **sigvec** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine.

---

## sigsuspend or sigpause Subroutine

### Purpose

Automatically changes the set of blocked signals and waits for a signal.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int sigsuspend ( SignalMask )
const sigset_t *SignalMask;
int sigpause ( SignalMask )
int SignalMask;
```

### Description

The **sigsuspend** subroutine replaces the signal mask of a thread with the set of signals pointed to by the *SignalMask* parameter. It then suspends execution of the thread until a signal is delivered that executes a signal-catching function or terminates the process. The **sigsuspend** subroutine does not allow the **SIGKILL** or **SIGSTOP** signal to be blocked. If a program attempts to block one of these signals, the **sigsuspend** subroutine gives no indication of the error.

If delivery of a signal causes the process to end, the **sigsuspend** subroutine does not return. If delivery of a signal causes a signal-catching function to start, the **sigsuspend** subroutine returns after the signal-catching function returns, with the signal mask restored to the set that existed prior to the **sigsuspend** subroutine.

The **sigsuspend** subroutine sets the signal mask and waits for an unblocked signal as one atomic operation. This means that signals cannot occur between the operations of setting the mask and waiting for a signal. If a program invokes the **sigprocmask (SIG\_SETMASK)** and **pause** subroutines separately, a signal that occurs between these subroutines might not be noticed by the **pause** subroutine.

In normal usage, a signal is blocked by using the **sigprocmask(SIG\_BLOCK,...)** subroutine for single-threaded applications, or the **sigthreadmask(SIG\_BLOCK,...)** subroutine for multi-threaded applications (using the **libpthread.a** threads library) at the beginning of a critical section. The process/thread then determines whether there is work for it to do. If no work is to be done, the process/thread waits for work by calling the **sigsuspend** subroutine with the mask previously returned by the **sigprocmask** or **sigthreadmask** subroutine.

The **sigpause** subroutine is provided for compatibility with older UNIX systems; its function is a subset of the **sigsuspend** subroutine.

## Parameter

*SignalMask*            Points to a set of signals.

## Return Values

If a signal is caught by the calling thread and control is returned from the signal handler, the calling thread resumes execution after the **sigsuspend** or **sigpause** subroutine, which always return a value of -1 and set the **errno** global variable to **EINTR**.

## Related Information

The **pause** subroutine, **sigprocmask** (“sigprocmask, sigsetmask, or sigblock Subroutine” on page 179) subroutine, **sigaction** or **signal** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine, **sigthreadmask** (“sigthreadmask Subroutine”) subroutine.

Signal Management in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs* provides more information about signal management in multi-threaded processes.

---

## sigthreadmask Subroutine

### Purpose

Sets the signal mask of a thread.

### Library

Threads Library (**libpthreads.a**)

### Syntax

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

int sigthreadmask( how, set, old_set)
int how;
const sigset_t *set;
sigset_t *old_set;
```

### Description

The **sigthreadmask** subroutine is used to examine or change the signal mask of the calling thread. The **sigprocmask** subroutine must not be used in a multi-threaded process.

Typically, the **sigthreadmask(SIG\_BLOCK)** subroutine is used to block signals during a critical section of code. The **sigthreadmask(SIG\_SETMASK)** subroutine is then used to restore the mask to the previous value returned by the **sigthreadmask(SIG\_BLOCK)** subroutine.

If there are any pending unblocked signals after the call to the **sigthreadmask** subroutine, at least one of those signals will be delivered before the **sigthreadmask** subroutine returns.

The **sigthreadmask** subroutine does not allow the **SIGKILL** or **SIGSTOP** signal to be blocked. If a program attempts to block either signal, the **sigthreadmask** subroutine gives no indication of the error.

**Note:** The `pthread.h` header file must be the first included file of each source file using the threads library.

## Parameters

<i>how</i>	Indicates the manner in which the set is changed. It can have one of the following values: <b>SIG_BLOCK</b> The resulting set is the union of the current set and the signal set pointed to by the <i>set</i> parameter. <b>SIG_UNBLOCK</b> The resulting set is the intersection of the current set and the complement of the signal set pointed to by the <i>set</i> parameter. <b>SIG_SETMASK</b> The resulting set is the signal set pointed to by the <i>set</i> parameter.
<i>set</i>	Specifies the signal set. If the value of the <i>Set</i> parameter is not null, it points to a set of signals to be used to change the currently blocked set. If the value of the <i>Set</i> parameter is null, the value of the <i>How</i> parameter is not significant and the process signal mask is unchanged. Thus, the call can be used to inquire about currently blocked signals.
<i>old_set</i>	If the <i>old_set</i> parameter is not the null value, the signal mask in effect at the time of the call is stored in the <i>spaced</i> pointed to by the <i>old_set</i> parameter.

## Return Values

Upon completion, a value of 0 is returned. If the `sigthreadmask` subroutine fails, the signal mask of the process is unchanged, a value of -1 is returned, and the global variable `errno` is set to indicate the error.

## Error Codes

The `sigthreadmask` subroutine is unsuccessful if the following is true:

<b>EFAULT</b>	The <i>set</i> or <i>old_set</i> pointers are not in the process address space.
<b>EINVAL</b>	The value of the <i>how</i> parameter is not supported.
<b>EPERM</b>	The calling thread does not have the privilege to change the signal's mask.

## Examples

To set the signal mask to block only the **SIGINT** signal from delivery, enter:

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

int return_value;
sigset_t newset;
sigset_t *newset_p;
. . .
newset_p = &newset;
sigemptyset(newset_p);
sigaddset(newset_p, SIGINT);
return_value = sigthreadmask(SIG_SETMASK, newset_p, NULL);
```

## Related Information

The `kill` or `killpg` subroutine, `pthread_kill` subroutine, `sigaction`, `sigvec`, or `signal` (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine, `sigpause` (“sigsuspend or sigpause Subroutine” on page 188) subroutine, `sigpending` (“sigpending Subroutine” on page 179) subroutine, `sigwait` (“sigwait Subroutine” on page 192) subroutine, `sigsuspend` (“sigsuspend or sigpause Subroutine” on page 188) subroutine.

## sigtimedwait and sigwaitinfo Subroutine

### Purpose

Waits for a signal, and provides a mechanism for retrieving any queued value.

### Library

Standard C Library (**libc.a**)

Threads Library (**libpthreads.a**)

### Syntax

```
#include <signal.h>

int sigtimedwait (set, info, timeout)
const sigset_t *set;
siginfo_t *info;
const struct timespec *timeout;

int sigwaitinfo (set, info)
const sigset_t *set;
siginfo_t *info;
```

### Description

The **sigwaitinfo** subroutine selects a pending signal from the set specified by the *set* parameter. If no signal in the *set* parameter is pending at the time of the call, the calling thread is suspended until one or more signals in the *set* parameter become pending or until it is interrupted by an unblocked, caught signal. If the wait was interrupted by an unblocked, caught signal, the subroutines will restart themselves.

The **sigwaitinfo** subroutine is functionally equivalent to the **sigwait** subroutine if the *info* argument is NULL. If the *info* argument is non-NULL, the **sigwaitinfo** subroutine is equivalent to the **sigwait** subroutine, except that the selected signal number is stored in the **si\_signo** member, and the cause of the signal is stored in the **si\_code** member of the *info* parameter. If any value is queued to the selected signal, the first such queued value is dequeued, and if the *info* argument is non-NULL, the value is stored in the **si\_value** member of the *info* parameter. If no further signals are queued for the selected signal, the pending indication for that signal is reset.

The **sigtimedwait** subroutine is equivalent to the **sigwaitinfo** subroutine except that if none of the signals specified by the *set* parameter are pending, the **sigtimedwait** subroutine waits for the time interval referenced by the *timeout* parameter. If the **timespec** structure pointed to by the *timeout* parameter contains a zero value and if none of the signals specified by the *set* parameter are pending, the **sigtimedwait** subroutine returns immediately with an error.

If there are multiple pending signals in the range **SIGRTMIN** to **SIGRTMAX**, the lowest numbered signal in that range will be selected.

**Note:** All signals in set should have been blocked prior to calling any of the **sigwait** subroutines.

### Parameters

<i>set</i>	Specifies the pending signals that may be selected.
<i>info</i>	Points to a <b>siginfo_t</b> in which additional signal information can be returned.

*timeout* Points to the **timespec** structure.

## Return Values

Upon successful completion, the **sigtimedwait** and **sigwaitinfo** subroutines return the selected signal number. If unsuccessful, the **sigtimedwait** and **sigwaitinfo** subroutines return -1 and set the **errno** variable to indicate the error.

## Error Codes

The **sigtimedwait** subroutine will fail if:

**EAGAIN** No signal specified by the *set* parameter was generated within the specified timeout period.

The **sigtimedwait** and **sigwaitinfo** subroutines may fail if:

**EINVAL** The *set* parameter is empty, or contains an invalid, non-catchable, or unsupported signal number.

The **sigtimedwait** subroutine may also fail when none of the selected signals are pending if:

**EINVAL** The *timeout* parameter specified a *tv\_nsec* value less than zero or greater than or equal to 1000 million.

## Related Information

“sigqueue Subroutine” on page 181 and “sigwait Subroutine”.

---

## sigwait Subroutine

### Purpose

Blocks the calling thread until a specified signal is received.

### Library

Threads Library (**libpthread.a**)

### Syntax

```
#include </usr/include/dce/cma_sigwait.h>
```

```
int sigwait ( set, sig)
const sigset_t *set;
int *sig;
```

### Description

The **sigwait** subroutine blocks the calling thread until one of the signal in the signal set *set* is received by the thread. Only asynchronous signals can be waited for.

The signal can be either sent directly to the thread, using the **pthread\_kill** subroutine, or to the process. In that case, the signal will be delivered to exactly one thread that has not blocked the signal.

Concurrent use of **sigaction** and **sigwait** subroutines on the same signal is forbidden.

## Parameters

*set* Specifies the set of signals to wait on.  
*sig* Points to where the received signal number will be stored.

## Return Values

Upon successful completion, the received signal number is returned via the *sig* parameter, and 0 is returned. Otherwise, an error code is returned.

## Error Code

The **sigwait** subroutine is unsuccessful if the following is true:

**EINVAL** The *set* parameter contains an invalid or unsupported signal number.

## Related Information

The **kill** subroutine, **pthread\_kill** subroutine, **sigaction** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine, **sigthreadmask** (“sigthreadmask Subroutine” on page 189) subroutine.

Signal Management in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs* .

---

## sin, sincf, or sinl Subroutine

### Purpose

Computes the sine.

### Syntax

```
#include <math.h>

double sin ( x )
double x;
float sincf ( x )
float x;
long double sinl ( x )
long double x;
```

### Description

The **sin**, **sincf**, **sinl** subroutines compute the sine of the *x* parameter, measured in radians.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

*x* Floating-point value  
*y* Floating-point value

## Return Values

Upon successful completion, the **sin**, **sinf**, and **sinl** subroutines return the sine of  $x$ .

If  $x$  is NaN, a NaN is returned.

If  $x$  is  $\pm 0$ ,  $x$  is returned.

If  $x$  is subnormal, a range error may occur and  $x$  should be returned.

If  $x$  is  $\pm\text{Inf}$ , a domain error occurs, and a NaN is returned.

## Error Codes

The **sin**, **sinf**, and **sinl** subroutines lose accuracy when passed a large value for the  $x$  parameter. In the **sin** subroutine, for example, values of  $x$  that are greater than  $\pi$  are argument-reduced by first dividing them by the machine value for  $2 * \pi$ , and then using the IEEE remainder of this division in place of  $x$ . Since the machine value of  $\pi$  can only approximate its infinitely precise value, the remainder of  $x/(2 * \pi)$  becomes less accurate as  $x$  becomes larger. Similar loss of accuracy occurs for the **sinl** subroutine during argument reduction of large arguments.

**sin** When the  $x$  parameter is extremely large, these functions return 0 when there would be a complete loss of significance. In this case, a message indicating **TLOSS** error is printed on the standard error output. For less extreme values causing partial loss of significance, a **PLOSS** error is generated but no message is printed. In both cases, the **errno** global variable is set to a **ERANGE** value.

These error-handling procedures may be changed with the **matherr** subroutine when using the **libmsaa.a** (**-lmsaa**) library.

## Related Information

The **matherr** subroutine, **sinh**, **sinhl** (“sinh, sinhf, or sinhl Subroutine”) subroutines.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

128-Bit long double Floating-Point Format in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

**math.h** in *AIX 5L Version 5.2 Files Reference*.

---

## sinh, sinhf, or sinhl Subroutine

### Purpose

Computes hyperbolic sine.

### Syntax

```
#include <math.h>
```

```
double sinh ( x )
```

```
double x;
```

```
float sinhf ( x )
```

```
float x;
```

```
long double sinhl (x)
double x;
```

## Description

The **sinh**, **sinhf**, and **sinhl** subroutines compute the hyperbolic sine of the *x* parameter.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies a double-precision floating-point value.

## Return Values

Upon successful completion, the **sinh**, **sinhf**, and **sinhl** subroutines return the hyperbolic sine of *x*.

If the result would cause an overflow, a range error occurs and **±HUGE\_VAL**, **±HUGE\_VALF**, and **±HUGE\_VALL** (with the same sign as *x*) is returned as appropriate for the type of the function.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$  or  $\pm \text{Inf}$ , *x* is returned.

If *x* is subnormal, a range error may occur and *x* should be returned.

## Error Codes

If the correct value overflows, the **sinh**, **sinhf** and **sinhl** subroutines return a correctly signed **HUGE\_VAL**, and the **errno** global variable is set to **ERANGE**.

These error-handling procedures should be changed with the **matherr** subroutine when the **libmsaa.a** (**-lmsaa**) library is used.

## Related Information

**asinh**, **acosh**, or **atanh** Subroutine, **feclearexcept** Subroutine, **fetestexcept** Subroutine, and **class**, **\_class**, **finite**, **isnan**, or **unordered** Subroutines in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

**math.h** in *AIX 5L Version 5.2 Files Reference*.

The **matherr** subroutine, **sin**, **asin**, **acos**, **atan**, or **atan2** (“**sin**, **sinf**, or **sinl** Subroutine” on page 193) subroutine.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

128-Bit long double Floating-Point Format in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## sleep, nsleep or usleep Subroutine

### Purpose

Suspends a current process from execution.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
unsigned int sleep ( Seconds)
```

```
#include <sys/time.h>
int nsleep ( Rqtp, Rmtp)
struct timestruc_t *Rqtp, *Rmtp;
```

```
int usleep ( Useconds)
useconds_t Useconds;
```

### Description

The **nsleep** subroutine is an extended form of the **sleep** subroutine. The **sleep** or **nsleep** subroutines suspend the current process until:

- The time interval specified by the *Rqtp* parameter elapses.
- A signal is delivered to the calling process that invokes a signal-catching function or terminates the process.
- The process is notified of an event through an event notification function.

The suspension time may be longer than requested due to the scheduling of other activity by the system. Upon return, the location specified by the *Rmtp* parameter shall be updated to contain the amount of time remaining in the interval, or 0 if the full interval has elapsed.

### Parameters

<i>Rqtp</i>	Time interval specified for suspension of execution.
<i>Rmtp</i>	Specifies the time remaining on the interval timer or 0.
<i>Seconds</i>	Specifies time interval in seconds.
<i>Useconds</i>	Specifies time interval in microseconds.

### Compatibility Interfaces

The **sleep** and **usleep** subroutines are provided to ensure compatibility with older versions of the operating system, AT&T System V and BSD systems. They are implemented simply as front-ends to the **nsleep** subroutine. Programs linking with the **libbsd.a** library get a BSD compatible version of the **sleep** subroutine. The return value from the BSD compatible **sleep** subroutine has no significance and should be ignored.

### Return Values

The **nsleep**, **sleep**, and **usleep** subroutines return a value of 0 if the requested time has elapsed.

If the **nsleep** subroutine returns a value of -1, the notification of a signal or event was received and the *Rmtp* parameter is updated to the requested time minus the time actually slept (unslept time), and the **errno** global variable is set.

If the **sleep** subroutine returns because of a premature arousal due to delivery of a signal, the return value will be the unslept amount (the requested time minus the time actually slept) in seconds.

## Error Codes

If the **nsleep** subroutine fails, a value of -1 is returned and the **errno** global variable is set to one of the following error codes:

- EINTR** A signal was caught by the calling process and control has been returned from the signal-catching routine, or the process has been notified of an event through an event notification function.
- EINVAL** The *Rqtp* parameter specified a nanosecond value less than zero or greater than or equal to one second.

The **sleep** subroutine is always successful and no return value is reserved to indicate an error.

## Related Information

The **alarm** subroutine, **pause** subroutine, **sigaction** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine.

List of Time Data Manipulation Services in *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

Subroutines Overview in *AIX 5L Version 5.2 System Management Guide: Operating System and Devices*.

---

## socketmark Subroutine

### Purpose

Determines whether a socket is at the out-of-band mark.

### Syntax

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

```
int socketmark(s)  
int s;
```

### Description

The **socketmark** subroutine determines whether the socket specified by the *s* parameter is at the out-of-band data mark. If the protocol for the socket supports out-of-band data by marking the stream with an out-of-band data mark, the **socketmark** subroutine returns a 1 when all data preceding the mark has been read and the out-of-band data mark is the first element in the receive queue. The **socketmark** subroutine does not remove the mark from the stream.

The use of this subroutine between receive operations allows an application to determine which received data precedes the out-of-band data and which follows the out-of-band data. There is an inherent race condition in the use of this function. On an empty receive queue, the current read of the location might well be at the mark, but the system has no way of knowing that the next data segment that will arrive from the network will carry the mark, and **socketmark** will return false. The next read operation will silently consume the mark. Because of this, the **socketmark** subroutine can only be used reliably when the application already knows that the out-of-band data has been seen by the system or that it is known that there is data waiting to be read at the socket.

## Parameters

*s* Specifies the socket to be checked.

## Return Values

Upon successful completion, the **socketmark** subroutine returns a value indicating whether the socket is at an out-of-band data mark. If the protocol has marked the data stream and all data preceding the mark has been read, the return value is 1. If there is no mark, or if data precedes the mark in the receive queue, the **socketmark** subroutine returns a 0. Otherwise, it returns a value of -1 and sets the **errno** global variable to indicate the error.

## Error Codes

**EBADF** The *s* parameter is not a valid file descriptor.  
**ENOTTY** The *s* parameter does not specify a descriptor for a socket.

---

## SpmiAddSetHot Subroutine

### Purpose

Adds a set of peer statistics values to a hotset.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
struct SpmiHotVals *SpmiAddSetHot(HotSet, StatName,
GrandParent, maxresp,
                                threshold, frequency, feed_type,
                                except_type, severity, trap_no)

struct SpmiHotSet *HotSet;
char *StatName;
SpmiCxHdl GrandParent;
int maxresp;
int threshold;
int frequency;
int feed_type;
int excp_type;
int severity;
int trap_no;
```

### Description

The **SpmiAddSetHot** subroutine adds a set of peer statistics to a hotset. The **SpmiHotSet** structure that provides the anchor point to the set must exist before the **SpmiAddSetHot** subroutine call can succeed.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### HotSet

Specifies a pointer to a valid structure of type **SpmiHotSet** as created by the **SpmiCreateHotSet** (“SpmiCreateHotSet” on page 201) subroutine call.

## StatName

Specifies the name of the statistic within the subcontexts (peer contexts) of the context identified by the *GrandParent* parameter.

## GrandParent

Specifies a valid **SpmiCxHdl** handle as obtained by another subroutine call. The handle must identify a context with at least one subcontext, which contains the statistic identified by the *StatName* parameter. If the context specified is one of the **RTime** contexts, no subcontext need to exist at the time the **SpmiAddSetHot** subroutine call is issued; the presence of the metric identified by the *StatName* parameter is checked against the context class description.

If the context specified has or may have multiple levels of instantiable context below it (such as the **FS** and **RTime/ARM** contexts), the metric is only searched for at the lowest context level. The **SpmiHotSet** created is a pseudo hotvals structure used to link together a peer group of **SpmiHotVals** structures, which are created under the covers, one for each subcontext of the *GrandParent* context. In the case of **RTime/ARM**, if additional contexts are later added under the *GrandParent* contexts, additional hotsets are added to the peer group. This is transparent to the application program, except that the **SpmiFirstHot**, **SpmiNextHot**, and **SpmiNextHotItem** subroutine calls will return the peer group **SpmiHotVals** pointer rather than the pointer to the pseudo structure.

Note that specifying a specific volume group context (such as **FS/rootvg**) or a specific application context (such as **RTime/ARN/armpeek**) is still valid and won't involve creation of pseudo **SpmiHotVals** structures.

## maxresp

Must be non-zero if *excp\_type* specifies that exceptions or SNMP traps must be generated. If specified as zero, indicates that all **SPMIHotItems** that meet the criteria specified by *threshold* must be returned, up-to a maximum of *maxresp* items. If both exceptions/traps and feeds are requested, the *maxresp* value is used to cap the number of exceptions/alerts as well as the number of items returned. If *feed\_type* is specified as **SiHotAlways**, the *maxresp* parameter is still used to return at most *maxresp* items.

Where the *GrandParent* argument specifies a context that has multiple levels of instantiable contexts below it, the *maxresp* is applied to each of the lowest level contexts above the the actual peer contexts at a time. For example, if the *GrandParent* context is **FS** (file systems) and the system has three volume groups, then a *maxresp* value of 2 could cause up to a maximum of  $2 \times 3 = 6$  responses to be generated.

## threshold

Must be non-zero if *excp\_type* specifies that exceptions or SNMP traps must be generated. If specified as zero, indicates that all values read qualify to be returned in feeds. The value specified is compared to the data value read for each peer statistic. If the data value exceeds the *threshold*, it qualifies to be returned as an **SpmiHotItems** element in the **SpmiHotVals** structure. If the *threshold* is specified as a negative value, the value qualifies if it is lower than the numeric value of *threshold*. If *feed\_type* is specified as **SiHotAlways**, the threshold value is ignored for feeds. For peer statistics of type **SiCounter**, the *threshold* must be specified as a rate per second; for **SiQuantity** statistics the *threshold* is specified as a level.

## frequency

Must be non-zero if *excp\_type* specifies that exceptions or SNMP traps must be generated. Ignored for feeds. Specifies the minimum number of minutes that must expire between any two exceptions/traps generated from this **SpmiHotVals** structure. This value must be specified as no less than 5 minutes.

## feed\_type

Specifies if feeds of **SpmiHotItems** should be returned for this **SpmiHotVals** structure. The following values are valid:

**SiHotNoFeed**

No feeds should be generated

**SiHotThreshold**

Feeds are controlled by *threshold*.

**SiHotAlways**

All values, up-to a maximum of *maxresp* must be returned as feeds.

**excp\_type**

Controls the generation of exception data packets and/or the generation of SNMP Traps from **xmservd**. Note that these types of packets and traps can only actually be sent if **xmservd** is running. Because of this, exception packets and SNMP traps are only generated as long as **xmservd** is active. Traps can only be generated on AIX systems. The conditions for generating exceptions and traps are controlled by the *threshold* and *frequency* parameters. The following values are valid for *excp\_type*:

**SiNoHotException**

Generate neither exceptions nor traps.

**SiHotException**

Generate exceptions but not traps.

**SiHotTrap**

Generate SNMP traps but not exceptions.

**SiHotBoth**

Generate both exceptions and SNMP traps.

**severity**

Required to be positive and greater than zero if exceptions are generated, otherwise specify as zero. Used to assign a severity code to the exception for display by **exmon**.

**trap\_no**

Required to be positive and greater than zero if SNMP traps are generated, otherwise specify as zero. Used to assign the trap number in the generated SNMP trap.

## Return Values

The **SpmiAddSetHot** subroutine returns a pointer to a structure of type **SpmiHotVals** if successful. If unsuccessful, the subroutine returns a NULL value.

## Programming Notes

The **SpmiAddSetHot** functions in a straight forward manner and as described previously in all cases where the *GrandParent* context is a context that has only one level of instantiable contexts below it. This covers most context types such as CPU, Disk, LAN, etc. In a few cases, currently only the **FS** (file system) and **RTime/ARM** (application response) contexts, the SPMI works by creating pseudo-hotvals structures that effectively expand the hotset. These pseudo-hotvals structures are created either at the time the **SpmiAddSetHot** call is issued or when new subcontexts are created for a context that's already the *GrandParent* of a hotvals peer set. For example:

When a peer set is created for **RTime/ARM**, maybe only a few or no subcontexts of this context exists. If two applications were defined at this point, say **checking** and **savings**, one valsset would be created for the **RTime/ARM** context and a pseudo-valset for each of **RTime/ARM/checking** and **RTime/ARM/savings**. As new applications are added to the **RTime/ARM** contexts, new pseudo-valsets are automatically added to the hotset.

Pseudo-valssets represent an implementation convenience and also helps minimize the impact of retrieving and presenting data for hotsets. As far as the caller of the **RSiGetHotItem** subroutine call is concerned, it is completely transparent. All this caller will ever see is the real hotvals structure. That is not the case for callers of **SpmiFirstHot**, **SpmiNextHot**, and **SpmiNextHotItem**. All of these subroutines will return pseudo-valssets and the calling program should be prepared to handle this.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

---

## SpmiCreateHotSet

### Purpose

Creates an empty hotset.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
struct SpmiHotSet *SpmiCreateHotSet()
```

### Description

The **SpmiCreateHotSet** subroutine creates an empty hotset and returns a pointer to an **SpmiHotSet** structure. This structure provides the anchor point for a hotset and must exist before the **SpmiAddSetHot** subroutine can be successfully called.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Return Values

The **SpmiCreateHotSet** subroutine returns a pointer to a structure of type **SpmiHotSet** if successful. If unsuccessful, the subroutine returns a NULL value.

### Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrMsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrMsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrMsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiDelSetHot Subroutine” on page 207
- “SpmiFreeHotSet Subroutine” on page 215
- “SpmiAddSetHot Subroutine” on page 198
- *Understanding SPMI Data Areas*

---

## SpmiCreateStatSet Subroutine

### Purpose

Creates an empty set of statistics.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
struct SpmiStatSet *SpmiCreateStatSet()
```

### Description

The **SpmiCreateStatSet** subroutine creates an empty set of statistics and returns a pointer to an **SpmiStatSet** structure.

The **SpmiStatSet** structure provides the anchor point to a set of statistics and must exist before the **SpmiPathAddSetStat** subroutine can be successfully called.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Return Values

The **SpmiCreateStatSet** subroutine returns a pointer to a structure of type **SpmiStatSet** if successful. If unsuccessful, the subroutine returns a NULL value.

### Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiDelSetStat Subroutine” on page 208
- “SpmiFreeStatSet Subroutine” on page 216
- “SpmiPathAddSetStat Subroutine” on page 234
- *Understanding SPMI Data Areas*

---

## SpmiDdsAddCx Subroutine

### Purpose

Adds a volatile context to the contexts defined by an application.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
char *SpmiDdsAddCx(Ix, Path, Descr, Asnno)
ushort Ix;
char *Path, *Descr;
int Asnno;
```

### Description

The **SpmiDdsAddCx** subroutine uses the shared memory area to inform the SPMI that a context is available to be added to the context hierarchy, moves a copy of the context to shared memory, and allocates memory for the data area.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Parameters

#### Ix

Specifies the element number of the added context in the table of dynamic contexts. No context can be added if the table of dynamic contexts has not been defined in the **SpmiDdsInit** subroutine call. The first element of the table is element number 0.

#### Path

Specifies the full path name of the context to be added. If the context is not at the top-level, the parent context must already exist.

#### *Descr*

Provides the description of the context to be added as it will be presented to data consumers.

#### **Asnno**

Specifies the ASN.1 number to be assigned to the new context. All subcontexts on the same level as the new context must have unique ASN.1 numbers. Typically, each time the **SpmiDdsAddCx** subroutine adds a subcontext to the same parent context, the **Asnno** parameter is incremented. See *Making Dynamic Data-Supplier Statistics Unique* for more information about ASN.1 numbers.

## **Return Values**

If successful, the **SpmiDdsAddCx** subroutine returns the address of the shared memory data area. If an error occurs, an error text is placed in the external **SpmiErrmsg** character array, and the subroutine returns a NULL value.

## **Error Codes**

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## **Files**

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## **Related Information**

For related information, see:

- “SpmiDdsDelCx Subroutine”
- “SpmiDdsInit Subroutine” on page 205

---

## **SpmiDdsDelCx Subroutine**

### **Purpose**

Deletes a volatile context.

### **Library**

SPMI Library (**libSpmi.a**)

### **Syntax**

```
#include sys/Spmidef.h
int SpmiDdsDelCx(Area)
char *Area;
```

## Description

The **SpmiDdsDelCx** subroutine informs the SPMI that a previously added, volatile context should be deleted.

If the SPMI has not detected that the context to delete was previously added dynamically, the **SpmiDdsDelCx** subroutine removes the context from the list of to-be-added contexts and returns the allocated shared memory to the free list. Otherwise, the **SpmiDdsDelCx** subroutine indicates to the SPMI that a context and its associated statistics must be removed from the context hierarchy and any allocated shared memory must be returned to the free list.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### Area

Specifies the address of the previously allocated shared memory data area as returned by an **SpmiDdsAddCx** subroutine call.

## Return Values

If successful, the **SpmiDdsDelCx** subroutine returns a value of 0. If an error occurs, an error text is placed in the external **SpmiErrmsg** character array, and the subroutine returns a nonzero value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiDdsAddCx Subroutine” on page 203
- “SpmiDdsInit Subroutine”
- *Understanding SPMI Data Areas*

---

## SpmiDdsInit Subroutine

### Purpose

- Establishes a program as a dynamic data-supplier (DDS) program.

## Library

SPMI Library (**libSpmi.a**)

## Syntax

```
#include sys/Spmidef.h
SpmiShare *SpmiDdsInit(CxTab, CxCnt, IxTab, IxCnt,
FileName)
cx_create *CxTab, *IxTab;
int CxCnt, IxCnt;
char *FileName;
```

## Description

The **SpmiDdsInit** subroutine establishes a program as a dynamic data-supplier (DDS) program. To do so, the **SpmiDdsInit** subroutine:

1. Determines the size of the shared memory required and creates a shared memory segment of that size.
2. Moves all static contexts and all statistics referenced by those contexts to the shared memory.
3. Calls the SPMI and requests it to add all of the DDS static contexts to the context tree.

### Notes:

1. The **SpmiDdsInit** subroutine issues an **SpmiInit** subroutine call if the application program has not issued one.
2. If the calling program uses shared memory for other purposes, including memory mapping of files, the **SpmiDdsInit** or the **SpmiInit** subroutine call must be issued before access is established to other shared memory areas.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### CxTab

Specifies a pointer to the table of nonvolatile contexts to be added.

### CxCnt

Specifies the number of elements in the table of nonvolatile contexts. Use the **CX\_L** macro to find this value.

### IxTab

Specifies a pointer to the table of volatile contexts the program may want to add later. If no contexts are defined, specify NULL.

### IxCnt

Specifies the number of elements in the table of volatile contexts. Use the **CX\_L** macro to find this value. If no contexts are defined, specify 0.

### FileName

Specifies the fully qualified path and file name to use when creating the shared memory segment. At execution time, if the file exists, the process running the DDS must be able to write to the file. Otherwise, the **SpmiDdsInit** subroutine call does not succeed. If the file does not exist, it is created. If the file cannot be created, the subroutine returns an error. If the file name includes directories that do not exist, the subroutine returns an error.

For non-AIX systems, a sixth argument is required to inform the SPMI how much memory to allocate in the DDS shared memory segment. This is not required for AIX systems because facilities exist to expand a memory allocation in shared memory. The sixth argument is:

#### size

Size in bytes of the shared memory area to allocate for the DDS program. This parameter is of type `int`.

## Return Values

If successful, the **SpmiDdsInit** subroutine returns the address of the shared memory control area. If an error occurs, an error text is placed in the external **SpmiErrMsg** character array, and the subroutine returns a NULL value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- `extern char SpmiErrMsg[];`
- `extern int SpmiErrno;`

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrMsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrMsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiExit Subroutine” on page 210
- “SpmiInit Subroutine” on page 223
- *Understanding SPMI Data Areas*

---

## SpmiDelSetHot Subroutine

### Purpose

Removes a single set of peer statistics from a hotset.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
int SpmiDelSetHot(HotSet, HotVal)
struct SpmiHotSet *HotSet;
struct SpmiHotVals *HotVal;
```

## Description

The **SpmiDelSetHot** subroutine removes a single set of peer statistics, identified by the *HotVal* parameter, from a hotset, identified by the *HotSet* parameter.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### HotSet

Specifies a pointer to a valid structure of type **SpmiHotSet**, as created by the “SpmiCreateHotSet” on page 201 subroutine call.

### HotVal

Specifies a pointer to a valid structure of type **SpmiHotVals**, as created by the “SpmiAddSetHot Subroutine” on page 198 subroutine call. You cannot specify an **SpmiHotVals** that was internally generated by the SPMI library code as described under the *GrandParent* parameter to **SpmiAddSetHot**.

## Return Values

The **SpmiDelSetHot** subroutine returns a value of 0 if successful. If unsuccessful, the subroutine returns a nonzero value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiCreateHotSet” on page 201
- “SpmiFreeHotSet Subroutine” on page 215
- “SpmiAddSetHot Subroutine” on page 198
- *Understanding SPMI Data Areas*

---

## SpmiDelSetStat Subroutine

### Purpose

Removes a single statistic from a set of statistics.

## Library

SPMI Library (**libSpmi.a**)

## Syntax

```
#include sys/Spmidef.h
int SpmiDelSetStat(StatSet, StatVal)
struct SpmiStatSet *StatSet;
struct SpmiStatVals *StatVal;
```

## Description

The **SpmiDelSetStat** subroutine removes a single statistic, identified by the *StatVal* parameter, from a set of statistics, identified by the *StatSet* parameter.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### StatSet

Specifies a pointer to a valid structure of type **SpmiStatSet** as created by the “SpmiCreateStatSet Subroutine” on page 202 subroutine call.

### StatVal

Specifies a pointer to a valid structure of type **SpmiStatVals** as created by the “SpmiPathAddSetStat Subroutine” on page 234 subroutine call.

## Return Values

The **SpmiDelSetStat** subroutine returns a value of 0 if successful. If unsuccessful, the subroutine returns a nonzero value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiCreateStatSet Subroutine” on page 202
- “SpmiFreeStatSet Subroutine” on page 216
- “SpmiPathAddSetStat Subroutine” on page 234

- *Understanding SPMI Data Areas*

---

## SpmiExit Subroutine

### Purpose

Terminates a dynamic data supplier (DDS) or local data consumer program's association with the SPMI, and releases allocated memory.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
void SpmiExit()
```

### Description

A successful "Spmilnit Subroutine" on page 223 or "SpmiDdslnit Subroutine" on page 205 call allocates shared memory. Therefore, a Dynamic Data Supplier (DDS) program that has issued a successful **Spmilnit** or **SpmiDdslnit** subroutine call should issue an **SpmiExit** subroutine call before the program exits the SPMI. Allocated memory is not released until the program issues an **SpmiExit** subroutine call.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

### Related Information

For related information, see:

- "Spmilnit Subroutine" on page 223
- "SpmiDdslnit Subroutine" on page 205

---

## SpmiFirstCx Subroutine

### Purpose

Locates the first subcontext of a context.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
struct SpmiCxLink *SpmiFirstCx(CxHandle)
SpmiCxHdl CxHandle;
```

## Description

The **SpmiFirstCx** subroutine locates the first subcontext of a context. The subroutine returns a NULL value if no subcontexts are found.

The structure pointed to by the returned pointer contains a handle to access the contents of the corresponding **SpmiCx** structure through the **SpmiGetCx** subroutine call.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### CxHandle

Specifies a valid **SpmiCxHdl** handle as obtained by another subroutine call.

## Return Values

The **SpmiFirstCx** subroutine returns a pointer to an **SpmiCxLink** structure if successful. If unsuccessful, the subroutine returns a NULL value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char **SpmiErrmsg**[];
- extern int **SpmiErrno**;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “**SpmiGetCx** Subroutine” on page 217
- “**SpmiNextCx** Subroutine” on page 226
- *Understanding SPMI Data Areas*
- *Understanding the SPMI Data Hierarchy*

---

## SpmiFirstHot Subroutine

### Purpose

Locates the first of the sets of peer statistics belonging to a hotset.

### Library

SPMI Library (**libSpmi.a**)

## Syntax

```
#include sys/Spmidef.h
struct SpmiHotVals *SpmiFirstHot(HotSet)
struct SpmiHotSet HotSet;
```

## Description

The **SpmiFirstHot** subroutine locates the first of the **SpmiHotVals** structures belonging to the specified **SpmiHotSet**. Using the returned pointer, the **SpmiHotSet** can then either be decoded directly by the calling program, or it can be used to specify the starting point for a subsequent **SpmiNextHotItem** subroutine call. The **SpmiFirstHot** subroutine should only be executed after a successful call to the **SpmiGetHotSet** subroutine.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### HotSet

Specifies a valid **SpmiHotSet** structure as obtained by another subroutine call.

## Return Values

The **SpmiFirstHot** subroutine returns a pointer to a structure of type **SpmiHotVals** structure if successful. If unsuccessful, the subroutine returns a NULL value. A returned pointer may refer to a pseudo-hotvals structure as described in the **SpmiAddSetHot** subroutine.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiCreateHotSet” on page 201
- “SpmiAddSetHot Subroutine” on page 198
- “SpmiNextHot Subroutine” on page 227
- “SpmiNextHotItem Subroutine” on page 228
- *Understanding SPMI Data Areas*
- *Understanding the SPMI Data Hierarchy*

---

## SpmiFirstStat Subroutine

### Purpose

Locates the first of the statistics belonging to a context.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
struct SpmiStatLink *SpmiFirstStat(CxHandle)
SpmiCxHdl CxHandle;
```

### Description

The **SpmiFirstStat** subroutine locates the first of the statistics belonging to a context. The subroutine returns a NULL value if no statistics are found.

The structure pointed to by the returned pointer contains a handle to access the contents of the corresponding **SpmiStat** structure through the “SpmiGetStat Subroutine” on page 219 call.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Parameters

#### CxHandle

Specifies a valid **SpmiCxHdl** handle as obtained by another subroutine call.

### Return Values

The **SpmiFirstStat** subroutine returns a pointer to a structure of type **SpmiStatLink** if successful. If unsuccessful, the subroutine returns a NULL value.

### Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

### Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiGetStat Subroutine” on page 219
- “SpmiNextStat Subroutine” on page 230
- *Understanding SPMI Data Areas*
- *Understanding the SPMI Data Hierarchy*

---

## SpmiFirstVals Subroutine

### Purpose

Returns a pointer to the first **SpmiStatVals** structure belonging to a set of statistics.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
struct SpmiStatVals *SpmiFirstVals(StatSet)
struct SpmiStatSet *StatSet;
```

### Description

The **SpmiFirstVals** subroutine returns a pointer to the first **SpmiStatVals** structure belonging to the set of statistics identified by the *StatSet* parameter. **SpmiStatVals** structures are accessed in reverse order so the last statistic added to the set of statistics is the first one returned. This subroutine call should only be issued after an **SpmiGetStatSet** subroutine has been issued against the statset.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Parameters

#### StatSet

Specifies a pointer to a valid structure of type **SpmiStatSet** as created by the **SpmiCreateStatSet** subroutine call.

### Return Values

The **SpmiFirstVals** subroutine returns a pointer to an **SpmiStatVals** structure if successful. If unsuccessful, the subroutine returns a NULL value.

### Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

`/usr/include/sys/Spmidef.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiCreateStatSet Subroutine” on page 202
- “SpmiNextVals Subroutine” on page 232
- *Understanding SPMI Data Areas*

---

## SpmiFreeHotSet Subroutine

### Purpose

Erases a hotset.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
int SpmiFreeHotSet(HotSet)
struct SpmiHotSet *HotSet;
```

### Description

The **SpmiFreeHotSet** subroutine erases the hotset identified by the *HotSet* parameter. All **SpmiHotVals** structures chained off the **SpmiHotSet** structure are deleted before the set itself is deleted.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Parameters

#### HotSet

Specifies a pointer to a valid structure of type **SpmiHotSet** as created by the “SpmiCreateHotSet” on page 201 subroutine call.

### Return Values

The **SpmiFreeHotSet** subroutine returns a value of 0 if successful. If unsuccessful, the subroutine returns a nonzero value.

### Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined

in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiCreateHotSet” on page 201
- “SpmiDelSetHot Subroutine” on page 207
- “SpmiAddSetHot Subroutine” on page 198
- *Understanding SPMI Data Areas*

---

## SpmiFreeStatSet Subroutine

### Purpose

Erases a set of statistics.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
int SpmiFreeStatSet(StatSet)
struct SpmiStatSet *StatSet;
```

### Description

The **SpmiFreeStatSet** subroutine erases the set of statistics identified by the *StatSet* parameter. All **SpmiStatVals** structures chained off the **SpmiStatSet** structure are deleted before the set itself is deleted.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Parameters

#### StatSet

Specifies a pointer to a valid structure of type **SpmiStatSet** as created by the **SpmiCreateStatSet** subroutine call.

### Return Values

The **SpmiFreeStatSet** subroutine returns a value of 0 if successful. If unsuccessful, the subroutine returns a nonzero value.

### Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];

- `extern int SpmiErrno;`

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h** Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiCreateStatSet Subroutine” on page 202
- “SpmiDelSetStat Subroutine” on page 208
- “SpmiPathAddSetStat Subroutine” on page 234
- *Understanding SPMI Data Areas*

---

## SpmiGetCx Subroutine

### Purpose

Returns a pointer to the **SpmiCx** structure corresponding to a specified context handle.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
struct SpmiCx *SpmiGetCx(CxHandle)
SpmiCxHdl CxHandle;
```

### Description

The **SpmiGetCx** subroutine returns a pointer to the **SpmiCx** structure corresponding to the context handle identified by the *CxHandle* parameter.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Parameters

#### CxHandle

Specifies a valid **SpmiCxHdl** handle as obtained by another subroutine call.

### Return Values

The **SpmiGetCx** subroutine returns a pointer to an **SpmiCx** data structure if successful. If unsuccessful, the subroutine returns NULL.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- `extern char SpmiErrmsg[];`
- `extern int SpmiErrno;`

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the `sys/Spmidef.h` file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

`/usr/include/sys/Spmidef.h` Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiFirstCx Subroutine” on page 210
- “SpmiNextCx Subroutine” on page 226
- *Understanding SPMI Data Areas*
- *Understanding the SPMI Data Hierarchy*

---

## SpmiGetHotSet Subroutine

### Purpose

Requests the SPMI to read the data values for all sets of peer statistics belonging to a specified **SpmiHotSet**.

### Library

SPMI Library (`libSpmi.a`)

### Syntax

```
#include sys/Spmidef.h
int SpmiGetHotSet(HotSet, Force);
struct SpmiHotSet *HotSet;
boolean Force;
```

### Description

The **SpmiGetHotSet** subroutine requests the SPMI to read the data values for all peer sets of statistics belonging to the **SpmiHotSet** identified by the *HotSet* parameter. The *Force* parameter is used to force the data values to be refreshed from their source.

The *Force* parameter works by resetting a switch held internally in the SPMI for all **SpmiStatVals** and **SpmiHotVals** structures, regardless of the **SpmiStatSets** and **SpmiHotSets** to which they belong. Whenever the data value for a peer statistic is requested, this switch is checked. If the switch is set, the SPMI reads the latest data value from the original data source. If the switch is not set, the SPMI reads the data value stored in the **SpmiHotVals** structure. This mechanism allows a program to synchronize and minimize the number of times values are retrieved from the source. One method programs can use is to ensure the force request is not issued more than once per elapsed amount of time.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### HotSet

Specifies a pointer to a valid structure of type **SpmiHotSet** as created by the “SpmiCreateHotSet” on page 201 subroutine call.

### Force

If set to true, forces a refresh from the original source before the SPMI reads the data values for the set. If set to false, causes the SPMI to read the data values as they were previously retrieved from the data source.

When the force argument is set true, the effect is that of marking all statistics known by the SPMI as obsolete, which causes the SPMI to refresh all requested statistics from kernel memory or other sources. As each statistic is refreshed, the obsolete mark is reset. Statistics that are not part of the **HotSet** specified in the subroutine call remain marked as obsolete. Therefore, if an application repetitively issues a series of, **SpmiGetHotSet** and **SpmiGetStatSet** subroutine calls for multiple hotsets and statsets, each time, only the first such call need set the force argument to true.

## Return Values

The **SpmiGetHotSet** subroutine returns a value of 0 if successful. If unsuccessful, the subroutine returns a nonzero value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiCreateHotSet” on page 201
- “SpmiAddSetHot Subroutine” on page 198

---

## SpmiGetStat Subroutine

### Purpose

Returns a pointer to the **SpmiStat** structure corresponding to a specified statistic handle.

## Library

SPMI Library (**libSpmi.a**)

## Syntax

```
#include sys/Spmidef.h
struct SpmiStat *SpmiGetStat(StatHandle)
SpmiStatHdl StatHandle;
```

## Description

The **SpmiGetStat** subroutine returns a pointer to the **SpmiStat** structure corresponding to the statistic handle identified by the *StatHandle* parameter.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### StatHandle

Specifies a valid **SpmiStatHdl** handle as obtained by another subroutine call.

## Return Values

The **SpmiGetStat** subroutine returns a pointer to a structure of type **SpmiStat** if successful. If unsuccessful, the subroutine returns a NULL value.

## Return Values

The **SpmiGetStat** subroutine returns a pointer to a structure of type **SpmiStat** if successful. If unsuccessful, the subroutine returns a NULL value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiFirstStat Subroutine” on page 213
- “SpmiNextStat Subroutine” on page 230
- *Understanding SPMI Data Areas*
- *Understanding the SPMI Data Hierarchy*

---

## SpmiGetStatSet Subroutine

### Purpose

Requests the SPMI to read the data values for all statistics belonging to a specified set.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
int SpmiGetStatSet(StatSet, Force);
struct SpmiStatSet *StatSet;
boolean Force;
```

### Description

The **SpmiGetStatSet** subroutine requests the SPMI to read the data values for all statistics belonging to the **SpmiStatSet** identified by the *StatSet* parameter. The *Force* parameter is used to force the data values to be refreshed from their source.

The *Force* parameter works by resetting a switch held internally in the SPMI for all **SpmiStatVals** and **SpmiHotVals** structures, regardless of the **SpmiStatSets** and **SpmiHotSets** to which they belong. Whenever the data value for a statistic is requested, this switch is checked. If the switch is set, the SPMI reads the latest data value from the original data source. If the switch is not set, the SPMI reads the data value stored for the **SpmiStatVals** structure. This mechanism allows a program to synchronize and minimize the number of times values are retrieved from the source. One method is to ensure the force request is not issued more than once per elapsed amount of time.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Parameters

#### StatSet

Specifies a pointer to a valid structure of type **SpmiStatSet** as created by the **SpmiCreateStatSet** subroutine call.

#### Force

If set to true, forces a refresh from the original source before the SPMI reads the data values for the set. If set to false, causes the SPMI to read the data values as they were previously retrieved from the data source.

When the force argument is set true, the effect is that of marking all statistics known by the SPMI as obsolete, which causes the SPMI to refresh all requested statistics from kernel memory or other sources. As each statistic is refreshed, the obsolete mark is reset. Statistics that are not part of the **StatSet** specified in the subroutine call remain marked as obsolete. Therefore, if an application repetitively issues the **SpmiGetStatSet** and **SpmiGetHotSet** subroutine calls for multiple statsets and hotsets, each time, only the first such call need set the force argument to true.

### Return Values

The **SpmiGetStatSet** subroutine returns a value of 0 if successful. If unsuccessful, the subroutine returns a nonzero value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- `extern char SpmiErrmsg[];`
- `extern int SpmiErrno;`

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

`/usr/include/sys/Spmidef.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiCreateStatSet Subroutine” on page 202
- “SpmiPathAddSetStat Subroutine” on page 234

---

## SpmiGetValue Subroutine

### Purpose

Returns a decoded value based on the type of data value extracted from the data field of an **SpmiStatVals** structure.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
float SpmiGetValue(StatSet, StatVal)
struct SpmiStatSet *StatSet;
struct SpmiStatVals *StatVal;
```

### Description

The **SpmiGetValue** subroutine performs the following steps:

1. Verifies that an **SpmiStatVals** structure exists in the set of statistics identified by the *StatSet* parameter.
2. Determines the format of the data field as being either **SiFloat** or **SiLong** and extracts the data value for further processing.
3. Determines the data value as being of either type **SiQuantity** or type **SiCounter**.
4. If the data value is of type **SiQuantity**, returns the **val** field of the **SpmiStatVals** structure.
5. If the data value is of type **SiCounter**, returns the value of the **val\_change** field of the **SpmiStatVals** structure divided by the elapsed number of seconds since the previous time a data value was requested for this set of statistics.

This subroutine call should only be issued after an **SpmiGetStatSet** subroutine has been issued against the statset.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### StatSet

Specifies a pointer to a valid structure of type **SpmiStatSet** as created by the **SpmiCreateStatSet** subroutine call.

### StatVal

Specifies a pointer to a valid structure of type **SpmiStatVals** as created by the **SpmiPathAddSetStat** subroutine call or returned by the **SpmiFirstVals** or **SpmiNextVals** subroutine calls.

## Return Values

The **SpmiGetValue** subroutine returns the decoded value if successful. If unsuccessful, the subroutine returns a negative value that has a numerical value of at least 1.1.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiGetStatSet Subroutine” on page 221
- “SpmiCreateStatSet Subroutine” on page 202
- “SpmiPathAddSetStat Subroutine” on page 234
- *Understanding SPMI Data Areas*

---

## Spmilnit Subroutine

### Purpose

Initializes the SPMI for a local data consumer program.

## Library

SPMI Library (**libSpmi.a**)

## Syntax

```
#include sys/Spmidef.h
int SpmiInit (TimeOut)
int TimeOut;
```

## Description

The **Spmilnit** subroutine initializes the SPMI. During SPMI initialization, a memory segment is allocated and the application program obtains basic addressability to that segment. An application program must issue the **Spmilnit** subroutine call before issuing any other subroutine calls to the SPMI.

**Note:** The **Spmilnit** subroutine is automatically issued by the **SpmiDdslnit** subroutine call. Successive **Spmilnit** subroutine calls are ignored.

**Note:** If the calling program uses shared memory for other purposes, including memory mapping of files, the **Spmilnit** subroutine call must be issued before access is established to other shared memory areas.

The SPMI entry point called by the **Spmilnit** subroutine assigns a segment register to be used by the SPMI subroutines (and the application program) for accessing common shared memory and establishes the access mode to the common shared memory segment. After SPMI initialization, the SPMI subroutines are able to access the common shared memory segment in read-only mode.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### TimeOut

Specifies the number of seconds the SPMI waits for a Dynamic Data Supplier (DDS) program to update its shared memory segment. If a DDS program does not update its shared memory segment in the time specified, the SPMI assumes that the DDS program has terminated or disconnected from shared memory and removes all contexts and statistics added by the DDS program.

The SPMI saves the largest *TimeOut* value received from the programs that invoke the SPMI. The *TimeOut* value must be zero or must be greater than or equal to 15 seconds and less than or equal to 600 seconds. A value of zero overrides any other value from any other program that invokes the SPMI and disables the checking for terminated DDS programs.

## Return Values

The **Spmilnit** subroutine returns a value of 0 if successful. If unsuccessful, the subroutine returns a nonzero value. If a nonzero value is returned, the application program should not attempt to issue additional SPMI subroutine calls.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiDdslnit Subroutine” on page 205
- “SpmiExit Subroutine” on page 210

---

## SpmiInstantiate Subroutine

### Purpose

Explicitly instantiates the subcontexts of an instantiable context.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
int SpmiInstantiate(CxHandle)
SpmiCxHdl CxHandle;
```

### Description

The **SpmiInstantiate** subroutine explicitly instantiates the subcontexts of an instantiable context. If the context is not instantiable, do not call the **SpmiInstantiate** subroutine.

An instantiation is done implicitly by the **SpmiPathGetCx** and **SpmiFirstCx** subroutine calls. Therefore, application programs usually do not need to instantiate explicitly.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Parameters

#### CxHandle

Specifies a valid context handle **SpmiCxHdl** as obtained by another subroutine call.

### Return Values

The **SpmiInstantiate** subroutine returns a value of 0 if successful. If the context is not instantiable, the subroutine returns a nonzero value.

### Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrMsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrMsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrMsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiFirstCx Subroutine” on page 210
- “SpmiPathGetCx Subroutine” on page 236
- *Understanding the SPMI Data Hierarchy*

---

## SpmiNextCx Subroutine

### Purpose

Locates the next subcontext of a context.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
struct SpmiCxLink *SpmiNextCx(CxLink )struct SpmiCxLink *CxLink;
```

### Description

The **SpmiNextCx** subroutine locates the next subcontext of a context, taking the context identified by the *CxLink* parameter as the current subcontext. The subroutine returns a NULL value if no further subcontexts are found.

The structure pointed to by the returned pointer contains an **SpmiCxHdl** handle to access the contents of the corresponding **SpmiCx** structure through the **SpmiGetCx** subroutine call.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Parameters

#### CxLink

Specifies a pointer to a valid **SpmiCxLink** structure as obtained by a previous **SpmiFirstCx** subroutine.

## Return Values

The **SpmiNextCx** subroutine returns a pointer to a structure of type **SpmiCxLink** if successful. If unsuccessful, the subroutine returns a NULL value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiFirstCx Subroutine” on page 210
- “SpmiGetCx Subroutine” on page 217
- *Understanding SPMI Data Areas*
- *Understanding the SPMI Data Hierarchy*

---

## SpmiNextHot Subroutine

### Purpose

Locates the next set of peer statistics **SpmiHotVals** belonging to an **SpmiHotSet**.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
struct SpmiHotVals *SpmiNextHot(HotSet, HotVals)
struct SpmiHotSet *HotSet;
struct SpmiHotVals *HotVals;
```

### Description

The **SpmiNextHot** subroutine locates the next **SpmiHotVals** structure belonging to an **SpmiHotSet**, taking the set of peer statistics identified by the *HotVals* parameter as the current one. The subroutine returns a NULL value if no further **SpmiHotVals** structures are found. The **SpmiNextHot** subroutine should only be executed after a successful call to the **SpmiGetHotSet** subroutine and (usually, but not necessarily) a call to the **SpmiFirstHot** subroutine and one or more subsequent calls to **SpmiNextHot**.

The subroutine allows the application programmer to position at the next set of peer statistics in preparation for using the **SpmiNextHotItem** subroutine call to traverse this peer set’s array of

**SpmiHotItems** elements. Use of this subroutine is only necessary if it is desired to skip over some **SpmiHotVals** structures in an **SpmiHotSet**. Under most circumstances, the **SpmiNextHotItem** will be the sole means of accessing all elements of the **SpmiHotItems** arrays of all peer sets belonging to an **SpmiHotSet**.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### HotSet

Specifies a valid pointer to an **SpmiHotSet** structure as obtained by a previous “SpmiCreateHotSet” on page 201 call.

### HotVals

Specifies a pointer to an **SpmiHotVals** structure as returned by a previous **SpmiFirstHot** or **SpmiNextHot** subroutine call or as returned by an **SpmiAddSetHot** subroutine call.

## Return Values

The **SpmiNextHot** subroutine returns a pointer to the next **SpmiHotVals** structure within the hotset. If no more **SpmiHotVals** structures are available, the subroutine returns a NULL value. A returned pointer may refer to a pseudo-hotvals structure as described the **SpmiAddSetHot** subroutine.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For more information, see:

- “SpmiFirstHot Subroutine” on page 211
- “SpmiGetHotSet Subroutine” on page 218
- “SpmiNextHotItem Subroutine”.
- *Data Access Structures and Handles, HotSets*

---

## SpmiNextHotItem Subroutine

### Purpose

Locates and decodes the next **SpmiHotItems** element at the current position in an **SpmiHotSet**.

## Library

SPMI Library (**libSpmi.a**)

## Syntax

```
#include sys/Spmidef.h
struct SpmiHotVals *SpmiNextHotItem(HotSet, HotVals, index,
value, name)
struct SpmiHotSet *HotSet;
struct SpmiHotVals *HotVals;
int *index;
float *value;
char **name;
```

## Description

The **SpmiNextHotItem** subroutine locates the next **SpmiHotItems** structure belonging to an **SpmiHotSet**, taking the element identified by the *HotVals* and *index* parameters as the current one. The subroutine returns a NULL value if no further **SpmiHotItems** structures are found. The **SpmiNextHotItem** subroutine should only be executed after a successful call to the **SpmiGetHotSet** subroutine.

The **SpmiNextHotItem** subroutine is designed to be used for walking all **SpmiHotItems** elements returned by a call to the **SpmiGetHotSet** subroutine, visiting the **SpmiHotVals** structures one by one. By feeding the returned value and the updated integer pointed to by *index* back to the next call, this can be done in a tight loop. Successful calls to **SpmiNextHotItem** will decode each **SpmiHotItems** element and return the data value in *value* and the name of the peer context that owns the corresponding statistic in *name*.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### HotSet

Specifies a valid pointer to an **SpmiHotSet** structure as obtained by a previous “SpmiCreateHotSet” on page 201 call.

### HotVals

Specifies a pointer to an **SpmiHotVals** structure as returned by a previous **SpmiNextHotItem**, **SpmiFirstHot**, or **SpmiNextHot** subroutine call or as returned by an **SpmiAddSetHot** subroutine call. If this parameter is specified as NULL, the first **SpmiHotVals** structure of the **SpmiHotSet** is used and the *index* parameter is assumed to be set to zero, regardless of its actual value.

### index

A pointer to an integer that contains the desired element number in the **SpmiHotItems** array of the **SpmiHotVals** structure specified by *HotVals*. A value of zero points to the first element. When the **SpmiNextHotItem** subroutine returns, the integer contain the index of the next **SpmiHotItems** element within the returned **SpmiHotVals** structure. If the last element of the array is decoded, the value in the integer will point beyond the end of the array, and the **SpmiHotVals** pointer returned will point to the peer set, which has now been completely decoded. By passing the returned **SpmiHotVals** pointer and the *index* parameter to the next call to **SpmiNextHotItem**, the subroutine will detect this and proceed to the first **SpmiHotItems** element of the next **SpmiHotVals** structure if one exists.

### value

A pointer to a float variable. A successful call will return the decoded data value for the statistic. Before the value is returned, the **SpmiNextHotItem** function:

- Determines the format of the data field as being either **SiFloat** or **SiLong** and extracts the data value for further processing.
- Determines the data value as being either type **SiQuantity** or type **SiCounter** and performs one of the actions listed here:
  - If the data value is of type **SiQuantity**, the subroutine returns the **val** field of the **SpmiHotItems** structure.
  - If the data value is of type **SiCounter**, the subroutine returns the value of the **val\_change** field of the **SpmiHotItems** structure divided by the elapsed number of seconds since the previous time a data value was requested for this set of statistics.

#### name

A pointer to a character pointer. A successful call will return a pointer to the name of the peer context for which the data value was read.

## Return Values

The **SpmiNextHotItem** subroutine returns a pointer to the current **SpmiHotVals** structure within the hotset. If no more **SpmiHotVals** structures are available, the subroutine returns a NULL value. The structure returned contains the data, such as threshold, which may be relevant for presentation of the results of an **SpmiGetHotSet** subroutine call to end-users. A returned pointer may refer to a pseudo-hotvals structure as described in the **SpmiAddSetHot** subroutine.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For more information, see:

- “SpmiFirstHot Subroutine” on page 211
- “SpmiNextHot Subroutine” on page 227
- “SpmiGetHotSet Subroutine” on page 218
- *Data Access Structures and Handles, HotSets*

---

## SpmiNextStat Subroutine

### Purpose

Locates the next statistic belonging to a context.

## Library

SPMI Library (**libSpmi.a**)

## Syntax

```
#include sys/Spmidef.h
struct SpmiStatLink *SpmiNextStat(StatLink)
struct SpmiStatLink *StatLink;
```

## Description

The **SpmiNextStat** subroutine locates the next statistic belonging to a context, taking the statistic identified by the *StatLink* parameter as the current statistic. The subroutine returns a NULL value if no further statistics are found.

The structure pointed to by the returned pointer contains an **SpmiStatHdl** handle to access the contents of the corresponding **SpmiStat** structure through the “SpmiGetStat Subroutine” on page 219 call.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### StatLink

Specifies a valid pointer to a **SpmiStatLink** structure as obtained by a previous “SpmiFirstStat Subroutine” on page 213 call.

## Return Values

The **SpmiNextStat** subroutine returns a pointer to a structure of type **SpmiStatLink** if successful. If unsuccessful, the subroutine returns a NULL value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiFirstStat Subroutine” on page 213
- “SpmiGetStat Subroutine” on page 219
- *Understanding SPMI Data Areas*

- *Understanding the SPMI Data Hierarchy*
- 

## SpmiNextVals Subroutine

### Purpose

Returns a pointer to the next **SpmiStatVals** structure in a set of statistics.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
struct SpmiStatVals *SpmiNextVals(StatSet, StatVal)
struct SpmiStatSet *StatSet;
struct SpmiStatVals *StatVal;
```

### Description

The **SpmiNextVals** subroutine returns a pointer to the next **SpmiStatVals** structure in a set of statistics, taking the structure identified by the *StatVal* parameter as the current structure. The **SpmiStatVals** structures are accessed in reverse order so the statistic added before the current one is returned. This subroutine call should only be issued after an **SpmiGetStatSet** subroutine has been issued against the statset.

### Parameters

#### StatSet

Specifies a pointer to a valid structure of type **SpmiStatSet** as created by the “SpmiCreateStatSet Subroutine” on page 202 call.

#### StatVal

Specifies a pointer to a valid structure of type **SpmiStatVals** as created by the “SpmiPathAddSetStat Subroutine” on page 234 subroutine call or returned by a previous “SpmiFirstVals Subroutine” on page 214 or **SpmiNextVals** subroutine call.

### Return Values

The **SpmiNextVals** subroutine returns a pointer to a **SpmiStatVals** structure if successful. If unsuccessful, the subroutine returns a NULL value.

---

## SpmiNextValue Subroutine

### Purpose

Returns either the first **SpmiStatVals** structure in a set of statistics or the next **SpmiStatVals** structure in a set of statistics and a decoded value based on the type of data value extracted from the data field of an **SpmiStatVals** structure.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h
```

```

struct SpmiStatVals*SpmiNextValue( StatSet, StatVal, value)
struct SpmiStatSet *StatSet;
struct SpmiStatVals *StatVal;
float *value;

```

## Description

Instead of issuing subroutine calls to “SpmiFirstVals Subroutine” on page 214 / “SpmiNextVals Subroutine” on page 232 (to get the first or next **SpmiStatVals** structure) followed by calls to **SpmiGetValue** (to get the decoded value from the **SpmiStatVals** structure), the **SpmiNextValue** subroutine returns both in one call. This subroutine call returns a pointer to the first **SpmiStatVals** structure belonging to the *StatSet* parameter if the *StatVal* parameter is NULL. If the *StatVal* parameter is not NULL, the next **SpmiStatVals** structure is returned, taking the structure identified by the *StatVal* parameter as the current structure. The data value corresponding to the returned **SpmiStatVals** structure is decoded and returned in the field pointed to by the value argument. In decoding the data value, the subroutine does the following:

- Determines the format of the data field as being either **SiFloat** or **SiLong** and extracts the data value for further processing.
- Determines the data value as being either type **SiQuantity** or type **SiCounter** and performs one of the actions listed here:
  - If the data value is of type **SiQuantity**, the subroutine returns the **val** field of the **SpmiStatVals** structure.
  - If the data value is of type **SiCounter**, the subroutine returns the value of the **val\_change** field of the **SpmiStatVals** structure divided by the elapsed number of seconds since the previous time a data value was requested for this set of statistics.

**Note:** This subroutine call should only be issued after an “SpmiGetStatSet Subroutine” on page 221 has been issued against the statset.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### StatSet

Specifies a pointer to a valid structure of type **SpmiStatSet** as created by the “SpmiCreateStatSet Subroutine” on page 202 call.

### StatVal

Specifies either a NULL pointer or a pointer to a valid structure of type **SpmiStatVals** as created by the “SpmiPathAddSetStat Subroutine” on page 234 call or returned by a previous **SpmiNextValue** subroutine call. If *StatVal* is NULL, then the first **SpmiStatVals** pointer belonging to the set of statistics pointed to by *StatSet* is returned.

**value**A pointer used to return a decoded value based on the type of data value extracted from the data field of the returned **SpmiStatVals** structure.

## Return Value

The **SpmiNextValue** subroutine returns a pointer to a **SpmiStatVals** structure if successful. If unsuccessful, the subroutine returns a NULL value.

If the **StatVal** parameter is:

### NULL

The first **SpmiStatVals** structure belonging to the **StatSet** parameter is returned.

## not NULL

The next **SpmiStatVals** structure after the structure identified by the **StatVal** parameter is returned and the value parameter is used to return a decoded value based on the type of data value extracted from the data field of the returned **SpmiStatVals** structure.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Programming Notes

The **SpmiNextValue** subroutine maintains internal state information so that retrieval of the next data value from a statset can be done without traversing linked lists of data structures. The stats information is kept separate for each process, but is shared by all threads of a process.

If the subroutine is accessed from multiple threads, the state information is useless and the performance advantage is lost. The same is true if the program is simultaneously accessing two or more statsets. To benefit from the performance advantage of the **SpmiNextValue** subroutine, a program should retrieve all values in order from one stat set before retrieving values from the next statset.

The implementation of the subroutine allows a program to retrieve data values beginning at any point in the statset if the **SpmiStatVals** pointer is known. Doing so will cause a linked list traversal. If subsequent invocations of **SpmiNextValue** uses the value returned from the first and following invocation as their second argument, the traversal of the link list can be avoided.

It should be noted that the value returned by a successful **SpmiNextValue** invocation is always the pointer to the **SpmiStatVals** structure whose data value is decoded and returned in the value argument.

## Files

<code>/usr/include/sys/Spmidef.h</code>	Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.
---	--

## Related Information

For related information, see:

- “SpmiGetStatSet Subroutine” on page 221
- “SpmiCreateStatSet Subroutine” on page 202
- “SpmiPathAddSetStat Subroutine”.
- *Data Access Structures and Handles, StatSets*

---

## SpmiPathAddSetStat Subroutine

### Purpose

Adds a statistics value to a set of statistics.

## Library

SPMI Library (**libSpmi.a**)

## Syntax

```
#include sys/Spmidef.h
struct SpmiStatVals *SpmiPathAddSetStat(StatSet, StatName,
Parent)
struct SpmiStatSet *StatSet;
char *StatName;
SpmiCxHdl Parent;
```

## Description

The **SpmiPathAddSetStat** subroutine adds a statistics value to a set of statistics. The **SpmiStatSet** structure that provides the anchor point to the set must exist before the **SpmiPathAddSetStat** subroutine call can succeed.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### StatSet

Specifies a pointer to a valid structure of type **SpmiStatSet** as created by the “SpmiCreateStatSet Subroutine” on page 202 call.

### StatName

Specifies the name of the statistic within the context identified by the *Parent* parameter. If the *Parent* parameter is NULL, you must specify the fully qualified path name of the statistic in the *StatName* parameter.

### Parent

Specifies either a valid **SpmiCxHdl** handle as obtained by another subroutine call or a NULL value.

## Return Values

The **SpmiPathAddSetStat** subroutine returns a pointer to a structure of type **SpmiStatVals** if successful. If unsuccessful, the subroutine returns a NULL value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- “SpmiGetStatSet Subroutine” on page 221
- “SpmiCreateStatSet Subroutine” on page 202
- “SpmiDelSetStat Subroutine” on page 208
- “SpmiFreeStatSet Subroutine” on page 216.
- *Data Access Structures and Handles, StatSets*

---

## SpmiPathGetCx Subroutine

### Purpose

Returns a handle to use when referencing a context.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h  
SpmiCxHdl SpmiPathGetCx(CxPath, Parent)  
char *CxPath;  
SpmiCxHdl Parent;
```

### Description

The **SpmiPathGetCx** subroutine searches the context hierarchy for a given path name of a context and returns a handle to use when subsequently referencing the context.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

### Parameters

#### CxPath

Specifies the path name of the context to find. If you specify the fully qualified path name in the *CxPath* parameter, you must set the *Parent* parameter to NULL. If the path name is not qualified or is only partly qualified (that is, if it does not include the names of all contexts higher in the data hierarchy), the **SpmiPathGetCx** subroutine begins searching the hierarchy at the context identified by the *Parent* parameter. If the *CxPath* parameter is either NULL or an empty string, the subroutine returns a handle identifying the Top context.

#### Parent

Specifies the anchor context that fully qualifies the *CxPath* parameter. If you specify a fully qualified path name in the *CxPath* parameter, you must set the *Parent* parameter to NULL.

### Return Values

The **SpmiPathGetCx** subroutine returns a handle to a context if successful. If unsuccessful, the subroutine returns a NULL value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- `extern char SpmiErrmsg[];`
- `extern int SpmiErrno;`

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

`/usr/include/sys/Spmidef.h`

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- *Understanding SPMI Data Areas*
- *Understanding the SPMI Data Hierarchy*

---

## SpmiStatGetPath Subroutine

### Purpose

Returns the full path name of a statistic.

### Library

SPMI Library (**libSpmi.a**)

### Syntax

```
#include sys/Spmidef.h>
char *miStatGetPath(Parent, StatHandle, MaxLevels)
SpmiCxHdlSp Parent;
SpmiStatHdl StatHandle;
int MaxLevels;
```

### Description

The **SpmiStatGetPath** subroutine returns the full path name of a statistic, given a parent context **SpmiCxHdl** handle and a statistics **SpmiStatHdl** handle. The *MaxLevels* parameter can limit the number of levels in the hierarchy that must be searched to generate the path name of the statistic.

The memory area pointed to by the returned pointer is freed when the **SpmiStatGetPath** subroutine call is repeated. For each invocation of the subroutine, a new memory area is allocated and its address returned. If the calling program needs the returned character string after issuing the **SpmiStatGetPath** subroutine call, the program must copy the returned string to locally allocated memory before reissuing the subroutine call.

This subroutine is part of the server option of the Performance Aide for AIX licensed product and is also included in the Performance Toolbox for AIX licensed product.

## Parameters

### Parent

Specifies a valid **SpmiCxHdl** handle as obtained by another subroutine call.

### StatHandle

Specifies a valid **SpmiStatHdl** handle as obtained by another subroutine call. This handle must point to a statistic belonging to the context identified by the *Parent* parameter.

### MaxLevels

Limits the number of levels in the hierarchy that must be searched to generate the path name. If this parameter is set to 0, no limit is imposed.

## Return Values

If successful, the **SpmiStatGetPath** subroutine returns a pointer to a character array containing the full path name of the statistic. If unsuccessful, the subroutine returns a NULL value.

## Error Codes

All SPMI subroutines use external variables to provide error information. To access these variables, an application program must define the following external variables:

- extern char SpmiErrmsg[];
- extern int SpmiErrno;

If the subroutine returns without an error, the **SpmiErrno** variable is set to 0 and the **SpmiErrmsg** character array is empty. If an error is detected, the **SpmiErrno** variable returns an error code, as defined in the **sys/Spmidef.h** file, and the **SpmiErrmsg** variable contains text, in English, explaining the cause of the error. See the *List of SPMI Error Codes* for more information.

## Files

**/usr/include/sys/Spmidef.h**

Declares the subroutines, data structures, handles, and macros that an application program can use to access the SPMI.

## Related Information

For related information, see:

- *Understanding SPMI Data Areas*
- *Understanding the SPMI Data Hierarchy*

---

## sqrt, sqrtf, or sqrtl Subroutine

### Purpose

Computes the square root.

### Syntax

```
#include <math.h>
double sqrt ( x)
double x;
float sqrtf (x)
float x;
```

```
long double sqrtl (x)
long double x;
```

## Description

The **sqrt**, **sqrtf**, and **sqrtl** subroutines compute the square root of the *x* parameter.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies some double-precision floating-point value.

## Return Values

Upon successful completion, the **sqrtf** subroutine returns the square root of *x*.

For finite values of  $x < -0$ , a domain error occurs, and a NaN is returned.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$  or  $+\text{Inf}$ , *x* is returned.

If *x* is  $-\text{Inf}$ , a domain error shall occur, and a NaN is returned.

## Error Codes

When using **libm.a (-lm)**:

For the **sqrt** subroutine, if the value of *x* is negative, a NaNQ is returned and the **errno** global variable is set to a **EDOM** value.

When using **libmsaa.a (-lmsaa)**:

If the value of *x* is negative, a 0 is returned and the **errno** global variable is set to a **EDOM** value. A message indicating a **DOMAIN** error is printed on the standard error output.

These error-handling procedures may be changed with the **matherr** subroutine when using the **libmsaa.a (-lmsaa)** library.

## Related Information

The **exp**, **expm1**, **log**, **log10**, **log1p**, or **pow** subroutine.

**feclearexcept** Subroutine, **fetestexcept** Subroutine, and **class**, **\_class**, **finite**, **isnan**, or **unordered** Subroutines in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

**math.h** in *AIX 5L Version 5.2 Files Reference*.

Subroutines Overview *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

## src\_err\_msg Subroutine

### Purpose

Retrieves a System Resource Controller (SRC) error message.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
int src_err_msg ( errno, ErrorText)
int errno;
char **ErrorText;
```

### Description

The **src\_err\_msg** subroutine retrieves a System Resource Controller (SRC) error message.

### Parameters

<i>errno</i>	Specifies the SRC error code.
<i>ErrorText</i>	Points to a character pointer to place the SRC error message.

### Return Values

Upon successful completion, the **src\_err\_msg** subroutine returns a value of 0. Otherwise, a value of -1 is returned. No error message is returned.

### Related Information

The **addssys** subroutine, **chssys** subroutine, **delssys** subroutine, **defssys** subroutine, **getsubsvr** subroutine, **getssys** subroutine, **srcsbuf** (“srcsbuf Subroutine” on page 244) subroutine, **srcrrqs** (“srcrrqs Subroutine” on page 241) subroutine, **srcsrpy** (“srcsrpy Subroutine” on page 250) subroutine, **srcsrqt** (“srcsrqt Subroutine” on page 253) subroutine, **srcstat** (“srcstat Subroutine” on page 259) subroutine, **srcstathdr** (“srcstathdr Subroutine” on page 264) subroutine, **srcstattxt** (“srcstattxt Subroutine” on page 265) subroutine, **srcstop** (“srcstop Subroutine” on page 266) subroutine, **srcstrt** (“srcstrt Subroutine” on page 269) subroutine.

List of SRC Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs.*

Programming Subsystem Communication with the SRC in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs.*

System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs.*

---

## src\_err\_msg\_r Subroutine

### Purpose

Gets the System Resource Controller (SRC) error message corresponding to the specified SRC error code.

### Library

System Resource Controller (**libsrc.a**)

### Syntax

```
#include <spc.h>
```

```
int src_err_msg_r (srcerrno, ErrorText)
int  srcerrno;
char ** ErrorText;
```

### Description

The **src\_err\_msg\_r** subroutine returns the message corresponding to the input `srcerrno` value in a caller-supplied buffer. This subroutine is threadsafe and reentrant.

### Parameters

<i>srcerrno</i>	Specifies the SRC error code.
<i>ErrorText</i>	Pointer to a variable containing the address of a caller-supplied buffer where the message will be returned. If the length of the message is unknown, the maximum message length can be used when allocating the buffer. The maximum message length is <b>SRC_BUF_MAX</b> in <b>/usr/include/spc.h</b> (2048 bytes).

### Return Values

Upon successful completion, the **src\_err\_msg\_r** subroutine returns a value of 0. Otherwise, no error message is returned and the subroutine returns a value of -1.

### Related Information

The **srcsbuf\_r** (“**srcsbuf\_r** Subroutine” on page 247), **srcsrqt\_r** (“**srcsrqt\_r** Subroutine” on page 256), **srcrrqs\_r** (“**srcrrqs\_r** Subroutine” on page 243), **srcstat\_r** (“**srcstat\_r** Subroutine” on page 262), and **srcstattxt\_r** (“**srcstattxt\_r** Subroutine” on page 265) subroutines.

List of SRC Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Programming Subsystem Communication with the SRC in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## srcrrqs Subroutine

### Purpose

Gets subsystem reply information from the System Resource Controller (SRC) request received.

## Library

System Resource Controller Library (**libsrc.a**)

## Syntax

```
#include <spc.h>
```

```
struct srchdr *srcrrqs ( Packet)  
char *Packet;
```

## Description

The **srcrrqs** subroutine saves the **srchdr** information contained in the packet the subsystem received from the System Resource Controller (SRC). The **srchdr** structure is defined in the **spc.h** file. This routine must be called by the subsystem to complete the reception process of any packet received from the SRC. The subsystem requires this information to reply to any request that the subsystem receives from the SRC.

**Note:** The saved **srchdr** information is overwritten each time this subroutine is called.

## Parameters

*Packet* Points to the SRC request packet received by the subsystem. If the subsystem received the packet on a message queue, the *Packet* parameter must point past the message type of the packet to the start of the request information. If the subsystem received the information on a socket, the *Packet* parameter points to the start of the packet received on the socket.

## Return Values

The **srcrrqs** subroutine returns a pointer to the static **srchdr** structure, which contains the return address for the subsystem response.

## Examples

The following will obtain the subsystem reply information:

```
int rc;  
struct sockaddr addr;  
int addrsz;  
struct srcreq packet;  
  
/* wait to receive packet from SRC daemon */  
rc=recvfrom(0, &packet, sizeof(packet), 0, &addr, &addrsz);  
/* grab the reply information from the SRC packet */  
if (rc>0)  
    srchdr=srcrrqs (&packet);
```

## Files

<b>/dev/SRC</b>	Specifies the <b>AF_UNIX</b> socket file.
<b>/dev/.SRC-unix</b>	Specifies the location for temporary socket files.

## Related Information

The **srcsbuf** (“srcsbuf Subroutine” on page 244) subroutine, **srcsrpy** (“srcsrpy Subroutine” on page 250) subroutine, **srcsrqt** (“srcsrqt Subroutine” on page 253) subroutine, **srcstat** (“srcstat Subroutine” on page 259) subroutine, **srcstathdr** (“srcstathdr Subroutine” on page 264) subroutine, **srcstattxt** (“srcstattxt Subroutine” on page 265) subroutine, **srcstop** (“srcstop Subroutine” on page 266) subroutine, **srcstrt** (“srcstrt Subroutine” on page 269) subroutine.

List of SRC Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Programming Subsystem Communication with the SRC in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## srcrrqs\_r Subroutine

### Purpose

Copies the System Resource Controller (SRC) request header to the specified buffer. The SRC request header contains the return address where the caller sends responses for this request.

### Library

System Resource Controller (**libsrc.a**)

### Syntax

```
#include <spc.h>
```

```
struct srchdr *srcrrqs_r (Packet, SRChdr)
char * Packet;
struct srchdr * SRChdr;
```

### Description

The **srcrrqs\_r** subroutine saves the SRC request header (srchdr) information contained in the packet the subsystem received from the Source Resource Controller. The **srchdr** structure is defined in the **spc.h** file. This routine must be called by the subsystem to complete the reception process of any packet received from the SRC. The subsystem requires this information to reply to any request that the subsystem receives from the SRC.

This subroutine is threadsafe and reentrant.

### Parameters

*Packet*        Points to the SRC request packet received by the subsystem. If the subsystem received the packet on a message queue, the *Packet* parameter must point past the message type of the packet to the start of the request information. If the subsystem received the information on a socket, the *Packet* parameter points to the start of the packet received on the socket.

*SRChdr*        Points to a caller-supplied buffer. The **srcrrqs\_r** subroutine copies the request header to this buffer.

### Examples

The following will obtain the subsystem reply information:

```
int rc;
struct sockaddr addr;
int addrsz;
struct srcreq packet;
struct srchdr *header;
struct srchdr *rtn_addr;

/*wait to receive packet from SRC daemon */
rc=recvfrom(0, &packet, sizeof(packet), 0, &addr, &addrsz;
```

```

/* grab the reply information from the SRC packet */
if (rc>0)
{
    header = (struct srchdr *)malloc(sizeof(struct srchdr));
    rtn_addr = srcrrqs_r(&packet,header);
    if (rtn_addr == NULL)
    {
        /* handle error */
        .
        .
    }
}

```

## Return Values

Upon successful completion, the **srcrrq\_r** subroutine returns the address of the caller-supplied buffer.

## Error Codes

If either of the input addresses is NULL, the **srcrrqs\_r** subroutine fails and returns a value of NULL.

**SRC\_PARM** One of the input addresses is NULL.

## Related Information

The **src\_err\_msg\_r** (“src\_err\_msg\_r Subroutine” on page 241), **srcsbuf\_r** (“srcsbuf\_r Subroutine” on page 247), **srcsrqt\_r** (“srcsrqt\_r Subroutine” on page 256), **srcstat\_r** (“srcstat\_r Subroutine” on page 262), and **srcstattxt\_r** (“srcstattxt\_r Subroutine” on page 265) subroutines.

---

## srcsbuf Subroutine

### Purpose

Gets status for a subserver or a subsystem and returns status text to be printed.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <spc.h>
```

```
intsrcsbuf(Host,Type,SubsystemName,
SubserverObject,SubsystemPID,StatusType,StatusFrom,StatusText,Continued)
```

```
char * Host, * SubsystemName;
```

```
char * SubserverObject, ** StatusText;
```

```
short Type, StatusType;
```

```
int SubsystemPID, StatusFrom, * Continued;
```

### Description

The **srcsbuf** subroutine gets the status of a subserver or subsystem and returns printable text for the status in the address pointed to by the *StatusText* parameter.

When the *StatusType* parameter is **SHORTSTAT** and the *Type* parameter is **SUBSYSTEM**, the **srcstat** subroutine is called to get the status of one or more subsystems. When the *StatusType* parameter is

**LONGSTAT** and the *Type* parameter is **SUBSYSTEM**, the **srcrsqt** subroutine is called to get the long status of one subsystem. When the *Type* parameter is not **SUBSYSTEM**, the **srcrsqt** subroutine is called to get the long or short status of a subserver.

## Parameters

<i>Host</i>	Specifies the foreign host on which this status action is requested. If the host is null, the status request is sent to the System Resource Controller (SRC) on the local host. The local user must be running as "root". The remote system must be configured to accept remote System Resource Controller requests. That is, the <b>srcmstr</b> daemon (see <b>/etc/inittab</b> ) must be started with the <b>-r</b> flag and the <b>/etc/hosts.equiv</b> or <b>.rhosts</b> file must be configured to allow remote requests.
<i>Type</i>	Specifies whether the status request applies to the subsystem or subserver. If the <i>Type</i> parameter is set to <b>SUBSYSTEM</b> , the status request is for a subsystem. If not, the status request is for a subserver and the <i>Type</i> parameter is a subserver code point.
<i>SubsystemName</i>	Specifies the name of the subsystem on which to get status. To get the status of all subsystems, use the <b>SRCALLSUBSYS</b> constant. To get the status of a group of subsystems, the <i>SubsystemName</i> parameter must start with the <b>SRCGROUP</b> constant, followed by the name of the group for which you want status appended. If you specify a null <i>SubsystemName</i> parameter, you must specify a <i>SubsystemPID</i> parameter.
<i>SubserverObject</i>	Specifies a subserver object. The <i>SubserverObject</i> parameter modifies the <i>Type</i> parameter. The <i>SubserverObject</i> parameter is ignored if the <i>Type</i> parameter is set to <b>SUBSYSTEM</b> . The use of the <i>SubserverObject</i> parameter is determined by the subsystem and the caller. This parameter will be placed in the <i>objname</i> field of the <b>subreq</b> structure that is passed to the subsystem.
<i>SubsystemPID</i>	Specifies the process ID of the subsystem on which to get status, as returned by the <b>srcstr</b> subroutine. You must specify the <i>SubsystemPID</i> parameter if multiple instances of the subsystem are active and you request a long subsystem status or subserver status. If you specify a null <i>SubsystemPID</i> parameter, you must specify a <i>SubsystemName</i> parameter.
<i>StatusType</i>	Specifies <b>LONGSTAT</b> for long status or <b>SHORTSTAT</b> for short status.
<i>StatusFrom</i>	Specifies whether status errors and messages are to be printed to standard output or just returned to the caller. When the <i>StatusFrom</i> parameter is <b>SSHHELL</b> , the errors are printed to standard output.
<i>StatusText</i>	Allocates memory for the printable text and sets the <i>StatusText</i> parameter to point to this memory. After it prints the text, the calling process must free the memory allocated for this buffer.
<i>Continued</i>	Specifies whether this call to the <b>srcsbuf</b> subroutine is a continuation of a status request. If the <i>Continued</i> parameter is set to <b>NEWREQUEST</b> , a request for status is sent and the <b>srcsbuf</b> subroutine then waits for another. On return, the <b>srcsbuf</b> subroutine is updated to the new continuation indicator from the reply packet and the <i>Continued</i> parameter is set to <b>END</b> or <b>STATCONTINUED</b> by the subsystem. If the <i>Continued</i> parameter is set to something other than <b>END</b> , this field must remain equal to that value; otherwise, this function will not be able to receive any more packets for the original status request. The calling process should not set the value of the <i>Continued</i> parameter to a value other than <b>NEWREQUEST</b> . The <i>Continued</i> parameter should not be changed while more responses are expected.

## Return Values

If the **srcsbuf** subroutine succeeds, it returns the size (in bytes) of printable text pointed to by the *StatusText* parameter.

## Error Codes

The **srcsbuf** subroutine fails if one or more of the following are true:

<b>SRC_BADSOCK</b>	The request could not be passed to the subsystem because of some socket failure.
--------------------	--

<b>SRC_CONT</b>	The subsystem uses signals. The request cannot complete.
<b>SRC_DMNA</b>	The SRC daemon is not active.
<b>SRC_INET_AUTHORIZED_HOST</b>	The local host is not in the remote <b>/etc/hosts.equiv</b> file.
<b>SRC_INET_INVALID_HOST</b>	On the remote host, the local host is not known.
<b>SRC_INVALID_USER</b>	The user is not root or group system.
<b>SRC_MMRV</b>	An SRC component could not allocate the memory it needs.
<b>SRC_NOCONTINUE</b>	The <i>Continued</i> parameter was not set to <b>NEWREQUEST</b> , and no continuation is currently active.
<b>SRC_NORPLY</b>	The request timed out waiting for a response.
<b>SRC_NSVR</b>	The subsystem is not active.
<b>SRC SOCK</b>	There is a problem with SRC socket communications.
<b>SRC_STPG</b>	The request was not passed to the subsystem. The subsystem is stopping.
<b>SRC_UDP</b>	The SRC port is not defined in the <b>/etc/services</b> file.
<b>SRC_UHOST</b>	The foreign host is not known.
<b>SRC_WICH</b>	There are multiple instances of the subsystem active.

## Examples

1. To get the status of a subsystem, enter:

```
char *status;
int continued=NEWREQUEST;
int rc;

do {
    rc=srctest("MaryC", SUBSYSTEM, "srctest", "", 0,
        SHORTSTAT, SSHELL, &status, continued);
    if (status!=0)
    {
        printf(status);
        free(status);
        status=0;
    }
} while (rc>0);
```

This gets short status of the srctest subsystem on the MaryC machine and prints the formatted status to standard output.

2. To get the status of a subserver, enter:

```
char *status;
int continued=NEWREQUEST;
int rc;

do {
    rc=srctest("", 12345, "srctest", "", 0,
        LONGSTAT, SSHELL, &status, continued);
    if (status!=0)
    {
        printf(status);
        free(status);
        status=0;
    }
} while (rc>0);
```

This gets long status for a specific subserver belonging to subsystem srctest. The subserver is the one having code point 12345. This request is processed on the local machine. The formatted status is printed to standard output.

## Files

<code>/etc/services</code>	Defines sockets and protocols used for Internet services.
<code>/dev/SRC</code>	Specifies the <b>AF_UNIX</b> socket file.
<code>/dev.SRC-unix</code>	Specifies the location for temporary socket files.

## Related Information

The **srcrrqs** (“srcrrqs Subroutine” on page 241) subroutine, **srcsrpy** (“srcsrpy Subroutine” on page 250) subroutine, **srcsrqt** (“srcsrqt Subroutine” on page 253) subroutine, **srcstat** (“srcstat Subroutine” on page 259) subroutine, **srcstathdr** (“srcstathdr Subroutine” on page 264) subroutine, **srcstattxt** (“srcstattxt Subroutine” on page 265) subroutine, **srcstop** (“srcstop Subroutine” on page 266) subroutine, **srcstrt** (“srcstrt Subroutine” on page 269) subroutine.

List of SRC Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Programming Subsystem Communication with the SRC in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## srcsbuf\_r Subroutine

### Purpose

Gets status for a subserver or a subsystem and returns status text to be printed.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <spc.h>
```

```
int srcsbuf_r(Host, Type, SubsystemName, SubserverObject, SubsystemPID,  
StatusType, StatusFrom, StatusText, Continued, SRCHandle)
```

```
char * Host, * SubsystemName;  
char * SubserverObject, ** StatusText;  
short Type, StatusType;  
pid_t SubsystemPID;  
int StatusFrom * Continued;  
char ** SRCHandle;
```

### Description

The **srcsbuf\_r** subroutine gets the status of a subserver or subsystem and returns printable text for the status in the address pointed to by the *StatusText* parameter. The **srcsbuf\_r** subroutine supports all the functions of the **srcbuf** subroutine except the *StatusFrom* parameter.

When the *StatusType* parameter is **SHORTSTAT** and the *Type* parameter is **SUBSYSTEM**, the **srcstat\_r** subroutine is called to get the status of one or more subsystems. When the *StatusType* parameter is

**LONGSTAT** and the *Type* parameter is **SUBSYSTEM**, the **srcsqt\_r** subroutine is called to get the long status of one subsystem. When the *Type* parameter is not **SUBSYSTEM**, the **srcsqt\_r** subroutine is called to get the long or short status of a subserver.

This routine is threadsafe and reentrant.

## Parameters

<i>Host</i>	Specifies the foreign host on which this status action is requested. If the host is null, the status request is sent to the System Resource Controller (SRC) on the local host.
<i>Type</i>	Specifies whether the status request applies to the subsystem or subserver. If the <i>Type</i> parameter is set to <b>SUBSYSTEM</b> , the status request is for a subsystem. If not, the status request is for a subserver and the <i>Type</i> parameter is a subserver code point.
<i>SubsystemName</i>	Specifies the name of the subsystem on which to get status. To get the status of all subsystems, use the <b>SRCALLSUBSYS</b> constant. To get the status of a group of subsystems, the <i>SubsystemName</i> parameter must start with the <b>SRCGROUP</b> constant, followed by the name of the group for which you want status appended. If you specify a null <i>SubsystemName</i> parameter, you must specify a <i>SubsystemPID</i> parameter.
<i>SubserverObject</i>	Specifies a subserver object. The <i>SubserverObject</i> parameter modifies the <i>Type</i> parameter. The <i>SubserverObject</i> parameter is ignored if the <i>Type</i> parameter is set to <b>SUBSYSTEM</b> . The use of the <i>SubserverObject</i> parameter is determined by the subsystem and the caller. This parameter will be placed in the objname field of the <b>subreq</b> structure that is passed to the subsystem.
<i>SubsystemPID</i>	Specifies the process ID of the subsystem on which to get status, as returned by the <b>srcstr</b> subroutine. You must specify the <i>SubsystemPID</i> parameter if multiple instances of the subsystem are active and you request a long subsystem status or subserver status. If you specify a null <i>SubsystemPID</i> parameter, you must specify a <i>SubsystemName</i> parameter.
<i>StatusType</i>	Specifies <b>LONGSTAT</b> for long status or <b>SHORTSTAT</b> for short status.
<i>StatusFrom</i>	Specifies whether status errors and messages are to be printed to standard output or just returned to the caller. When the <i>StatusFrom</i> parameter is <b>SSHELL</b> , the errors are printed to standard output. The <b>SSHELL</b> value is not recommended in a multithreaded environment since error messages to standard output may be interleaved in an unexpected manner.
<i>StatusText</i>	Allocates memory for the printable text and sets the <i>StatusText</i> parameter to point to this memory. After it prints the text, the calling process must free the memory allocated for this buffer.
<i>Continued</i>	Specifies whether this call to the <b>srcsbuf_r</b> subroutine is a continuation of a status request. If the <i>Continued</i> parameter is set to <b>NEWREQUEST</b> , a request for status is sent and the <b>srcsbuf_r</b> subroutine then waits for a reply. On return from the <b>srcsbuf_r</b> subroutine, the <i>Continued</i> parameter is updated to the new continuation indicator from the reply packet. The continuation indicator in the reply packet will be set to <b>END</b> or <b>STATCONTINUED</b> by the subsystem. If the <i>Continued</i> parameter is set to something other than <b>END</b> , the caller should not change that value; otherwise, this function will not be able to receive any more packets for the original status request. The calling process should not set the value of the <i>Continued</i> parameter to a value other than <b>NEWREQUEST</b> . In normal processing, the <i>Continued</i> parameter should not be changed while more responses are expected. The caller must continue to call the <b>srcsbuf_r</b> subroutine until <b>END</b> is received. As an alternative, call the <b>srcsbuf_r</b> subroutine with <b>Continued=SRC_CLOSE</b> to discard the remaining data, close the socket, and free the internal buffers.
<i>SRCHandle</i>	Identifies a request and its associated responses. Set to NULL by the caller for a <b>NEWREQUEST</b> . The <b>srcsbuf_r</b> subroutine saves a value in <i>SRCHandle</i> to allow <b>srcsbuf_r</b> continuation calls to use the same socket and internal buffers. The <i>SRCHandle</i> parameter should not be changed by the caller except for <b>NEWREQUESTs</b> .

## Return Values

If the `srcsbuf_r` subroutine succeeds, it returns the size (in bytes) of printable text pointed to by the `StatusText` parameter.

## Error Codes

The `srcsbuf_r` subroutine fails and returns the corresponding error code if one of the following error conditions is detected:

<b>SRC_BADSOCK</b>	The request could not be passed to the subsystem because of some socket failure.
<b>SRC_CONT</b>	The subsystem uses signals. The request cannot complete.
<b>SRC_DMNA</b>	The SRC daemon is not active.
<b>SRC_INET_AUTHORIZED_HOST</b>	The local host is not in the remote <code>/etc/hosts.equiv</code> file.
<b>SRC_INET_INVALID_HOST</b>	On the remote host, the local host is not known.
<b>SRC_INVALID_USER</b>	The user is not root or group system.
<b>SRC_MMRV</b>	An SRC component could not allocate the memory it needs.
<b>SRC_NOCONTINUE</b>	The <code>Continued</code> parameter was not set to <b>NEWREQUEST</b> , and no continuation is currently active.
<b>SRC_NORPLY</b>	The request timed out waiting for a response.
<b>SRC_NSVR</b>	The subsystem is not active.
<b>SRC_SOCKET</b>	There is a problem with SRC socket communications.
<b>SRC_STPG</b>	The request was not passed to the subsystem. The subsystem is stopping.
<b>SRC_UDP</b>	The SRC port is not defined in the <code>/etc/services</code> file.
<b>SRC_UHOST</b>	The foreign host is not known.
<b>SRC_WICH</b>	There are multiple instances of the subsystem active.

## Examples

1. To get the status of a subsystem, enter:

```
char *status;
int continued=NEWREQUEST;
int rc;
char *handle

do {
    rc=srcsbuf_r("MaryC", SUBSYSTEM, "srctest", "", 0,
        SHORTSTAT, SDAEMON, &status, continued, &handle);
    if (status!=0)
    {
        printf(status);
        free(status);
        status=0;
    }
} while (rc>0);
if (rc<0)
{
    ...handle error from srcsbuf_r...
}
```

This gets short status of the `srctest` subsystem on the `MaryC` machine and prints the formatted status to standard output.

**Caution:** In a multithreaded environment, the caller must manage the sharing of standard output between threads. Set the `StatusFrom` parameter to `SDAEMON` to prevent unexpected error messages from being printed to standard output.

2. To get the status of a subserver, enter:

```

char *status;
int continued=NEWREQUEST;
int rc;
char *handle

do {
    rc=srcsbuf_r("", 12345, "srctest", "", 0,
        LONGSTAT, SDAEMON, &status, continued, &handle);
    if (status!=0)
        {
            printf(status);
            free(status);
            status=0;
        }
} while (rc>0);
if (rc<0)
{
    ...handle error from srcsbuf_r...
}

```

This gets long status for a specific subserver belonging to subsystem srctest. The subserver is the one having code point 12345. This request is processed on the local machine. The formatted status is printed to standard output.

#### **CAUTION:**

**In a multithreaded environment, the caller must manage the sharing of standard output between threads. Set the *StatusFrom* parameter to SDAEMON to prevent unexpected error messages from being printed to standard output.**

## **Related Information**

The **src\_err\_msg\_r** (“src\_err\_msg\_r Subroutine” on page 241) subroutine, **srcsrqt\_r** (“srcsrqt\_r Subroutine” on page 256) subroutine, **srcrrqs\_r** (“srcrrqs\_r Subroutine” on page 243) subroutine, **srcstat\_r** (“srcstat\_r Subroutine” on page 262) subroutine, **srcstattxt\_r** (“srcstattxt\_r Subroutine” on page 265) subroutine.

List of SRC Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Programming Subsystem Communication with the SRC in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## **srcsrpy Subroutine**

### **Purpose**

Sends a reply to a request from the System Resource Controller (SRC) back to the client process.

### **Library**

System Resource Controller Library (**libsrc.a**)

### **Syntax**

```
#include <spc.h>
```

```
int srcsrpy ( SRChdr, PPacket, PPacketSize, Continued)
```

```

struct srchdr *SRChdr;
char *Ppck;
int PpckSize;
ushort Cont;

```

## Description

The **srcsrpy** subroutine returns a subsystem reply to a System Resource Controller (SRC) subsystem request. The format and content of the reply are determined by the subsystem and the requester, but must start with a **srchdr** structure. This structure and all others required for subsystem communication with the SRC are defined in the **/usr/include/spc.h** file. The subsystem must reply with a pre-defined format and content for the following requests: **START**, **STOP**, **STATUS**, **REFRESH**, and **TRACE**. The **START**, **STOP**, **REFRESH**, and **TRACE** requests must be answered with a **srcrep** structure. The **STATUS** request must be answered with a reply in the form of a **statbuf** structure.

**Note:** The **srcsrpy** subroutine creates its own socket to send the subsystem reply packets.

## Parameters

<i>SRChdr</i>	Points to the reply address buffer as returned by the <b>srcrrqs</b> subroutine.						
<i>Ppck</i>	Points to the reply packet. The first element of the reply packet is a <b>srchdr</b> structure. The cont element of the <i>Ppck</i> -> <b>srchdr</b> structure is modified on returning from the <b>srcsrpy</b> subroutine. The second element of the reply packet should be a <b>svrreply</b> structure, an array of <b>statcode</b> structures, or another format upon which the subsystem and the requester have agreed.						
<i>PpckSize</i>	Specifies the number of bytes in the reply packet pointed to by the <i>Ppck</i> parameter. The <i>PpckSize</i> parameter may be the size of a <b>short</b> , or it may be between the size of a <b>srchdr</b> structure and the <b>SRCPKTMAX</b> value, which is defined in the <b>spc.h</b> file.						
<i>Cont</i>	Indicates whether this reply is to be continued. If the <i>Cont</i> parameter is set to the constant <b>END</b> , no more reply packets are sent for this request. If the <i>Cont</i> parameter is set to <b>CONTINUED</b> , the second element of what is indicated by the <i>Ppck</i> parameter must be a <b>svrreply</b> structure, since the <i>rtmsg</i> element of the <b>svrreply</b> structure is printed to standard output. For a status reply, the <i>Cont</i> parameter is set to <b>STATCONTINUED</b> , and the second element of what is pointed to by the <i>Ppck</i> parameter must be an array of <b>statcode</b> structures. If a STOP subsystem request is received, only one reply packet can be sent and the <i>Cont</i> parameter must be set to <b>END</b> . Other types of continuations, as determined by the subsystem and the requester, must be defined using positive values for the <i>Cont</i> parameter. Values other than the following must be used:						
	<table> <tr> <td>0</td> <td>END</td> </tr> <tr> <td>1</td> <td>CONTINUED</td> </tr> <tr> <td>2</td> <td>STATCONTINUED</td> </tr> </table>	0	END	1	CONTINUED	2	STATCONTINUED
0	END						
1	CONTINUED						
2	STATCONTINUED						

## Return Values

If the **srcsrpy** subroutine succeeds, it returns the value **SRC\_OK**.

## Error Codes

The **srcsrpy** subroutine fails if one or both of the following are true:

<b>SRC_SOCKET</b>	There is a problem with SRC socket communications.
<b>SRC_REPLYSZ</b>	SRC reply size is invalid.

## Examples

1. To send a **STOP** subsystem reply, enter:

```

struct srcrep return_packet;
struct srchdr *srchdr;

bzero(&return_packet,sizeof(return_packet));
return_packet.svrreply.rtncode=SRC_OK;
strcpy(return_packet.svrreply,"srctest");

srcsrpy(srchdr,return_packet,sizeof(return_packet),END);

```

This entry sends a message that the subsystem srctest is stopping successfully.

2. To send a **START** subserver reply, enter:

```

struct srcrep return_packet;
struct srchdr *srchdr;

bzero(&return_packet,sizeof(return_packet));
return_packet.svrreply.rtncode=SRC_SUBMSG;
strcpy(return_packet.svrreply,objname,"mysubserver");
strcpy(return_packet.svrreply,objtext,"The subserver,\
mysubserver, has been started");

srcsrpy(srchdr,return_packet,sizeof(return_packet),END);

```

The resulting message indicates that the start subserver request was successful.

3. To send a status reply, enter:

```

int rc;
struct sockaddr addr;
int addrsz;
struct srcreq packet;
struct
{
    struct srchdr srchdr;
    struct statcode statcode[10];
} status;
struct srchdr *srchdr;
struct srcreq packet;
.
.
.
/* grab the reply information from the SRC packet */
srchdr=srcrrqs(&packet);
bzero(&status.statcode[0].objname,

/* get SRC status header */
srcstathdr(status.statcode[0].objname,
    status.statcode[0].objtext);
.
.
.
/* send status packet(s) */
srcsrpy(srchdr,&status,sizeof(status),STATCONTINUED);
.
.
.
srcsrpy(srchdr,&status,sizeof(status),STATCONTINUED);

/* send final packet */
srcsrpy(srchdr,&status,sizeof(struct srchdr),END);

```

This entry sends several status packets.

## Files

**/dev.SRC-unix** Specifies the location for temporary socket files.

## Related Information

The **srcrrqs** (“srcrrqs Subroutine” on page 241) subroutine, **srcsbuf** (“srcsbuf Subroutine” on page 244) subroutine, **srcsrqt** (“srcsrqt Subroutine”) subroutine, **srcstat** (“srcstat Subroutine” on page 259) subroutine, **srcstathdr** (“srcstathdr Subroutine” on page 264) subroutine, **srcstattxt** (“srcstattxt Subroutine” on page 265) subroutine, **srcstop** (“srcstop Subroutine” on page 266) subroutine, **srcstrt** (“srcstrt Subroutine” on page 269) subroutine.

List of SRC Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Programming Subsystem Communication with the SRC in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding SRC Communication Types in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## srcsrqt Subroutine

### Purpose

Sends a request to a subsystem.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <spc.h> srcsrqt(Host, SubsystemName, SubsystemPID,  
RequestLength, SubsystemRequest, ReplyLength, ReplyBuffer, StartItAlso, Continued)
```

```
char * Host, * SubsystemName;
```

```
char * SubsystemRequest, * ReplyBuffer;
```

```
int SubsystemPID, StartItAlso, * Continued;
```

```
short RequestLength, * ReplyLength;
```

### Description

The **srcsrqt** subroutine sends a request to a subsystem, waits for a response, and returns one or more replies to the caller. The format of the request and the reply is determined by the caller and the subsystem.

**Note:** The **srcsrqt** subroutine creates its own socket to send a request to the subsystem. The socket that this function opens remains open until an error or an end packet is received.

Two types of continuation are returned by the **srcsrqt** subroutine:

No continuation                      *ReplyBuffer->srchdr.continued* is set to the **END** constant.

Reply continuation

*ReplyBuffer*->*srchrdr.continued* is not set to the **END** constant, but to a positive value agreed upon by the calling process and the subsystem. The packet is returned to the caller.

## Parameters

<i>SubsystemPID</i>	The process ID of the subsystem.
<i>Host</i>	Specifies the foreign host on which this subsystem request is to be sent. If the host is null, the request is sent to the subsystem on the local host. The local user must be running as "root". The remote system must be configured to accept remote System Resource Controller requests. That is, the <b>srcmstr</b> daemon (see <i>/etc/inittab</i> ) must be started with the <b>-r</b> flag and the <i>/etc/hosts.equiv</i> or <i>.rhosts</i> file must be configured to allow remote requests.
<i>SubsystemName</i>	Specifies the name of the subsystem to which this request is to be sent. You must specify a <i>SubsystemName</i> if you do not specify a <i>SubsystemPID</i> .
<i>RequestLength</i>	Specifies the length, in bytes, of the request to be sent to the subsystem. The maximum value in bytes for this parameter is 2000 bytes.
<i>SubsystemRequest</i>	Points to the subsystem request packet.
<i>ReplyLength</i>	Specifies the maximum length, in bytes, of the reply to be received from the subsystem. On return from the <b>srcsrqt</b> subroutine, the <i>ReplyLength</i> parameter is set to the actual length of the subsystem reply packet.
<i>ReplyBuffer</i>	Points to a buffer for the receipt of the reply packet from the subsystem.
<i>StartItAlso</i>	Specifies whether the subsystem should be started if it is nonactive. When nonzero, the System Resource Controller (SRC) attempts to start a nonactive subsystem, and then passes the request to the subsystem.
<i>Continued</i>	Specifies whether this call to the <b>srcsrqt</b> subroutine is a continuation of a previous request. If the <i>Continued</i> parameter is set to <b>NEWREQUEST</b> , a request for it is sent to the subsystem and the subsystem is notified that another response is expected. The calling process should never set <i>Continued</i> to any value other than <b>NEWREQUEST</b> . The last response from the subsystem will set <i>Continued</i> to <b>END</b> .

## Return Values

If the **srcsrqt** subroutine is successful, the value **SRC\_OK** is returned.

## Error Codes

The **srcsrqt** subroutine fails if one or more of the following are true:

<b>SRC_BADSOCK</b>	The request could not be passed to the subsystem because of a socket failure.
<b>SRC_CONT</b>	The subsystem uses signals. The request cannot complete.
<b>SRC_DMNA</b>	The SRC daemon is not active.
<b>SRC_INET_AUTHORIZED_HOST</b>	The local host is not in the remote <i>/etc/hosts.equiv</i> file.
<b>SRC_INET_INVALID_HOST</b>	On the remote host, the local host is not known.
<b>SRC_INVALID_USER</b>	The user is not root or group system.
<b>SRC_MMRV</b>	An SRC component could not allocate the memory it needs.
<b>SRC_NOCONTINUE</b>	The <i>Continued</i> parameter was not set to <b>NEWREQUEST</b> , and no continuation is currently active.
<b>SRC_NORPLY</b>	The request timed out waiting for a response.
<b>SRC_NSVR</b>	The subsystem is not active.
<b>SRC_REQLEN2BIG</b>	The <i>RequestLength</i> is greater than the maximum 2000 bytes.
<b>SRC_SOCK</b>	There is a problem with SRC socket communications.
<b>SRC_STPG</b>	The request was not passed to the subsystem. The subsystem is stopping.
<b>SRC_UDP</b>	The SRC port is not defined in the <i>/etc/services</i> file.
<b>SRC_UHOST</b>	The foreign host is not known.

## Examples

1. To request long subsystem status, enter:

```
int cont=NEWREQUEST;
int rc;
short replen;
short reqlen;
struct
{
    struct srchdr srchdr;
    struct statcode statcode[20];
} statbuf;
struct subreq subreq;

subreq.action=STATUS;
subreq.object=SUBSYSTEM;
subreq.parm1=LONGSTAT;
strcpy(subreq.objname,"srctest");
replen=sizeof(statbuf);
reqlen=sizeof(subreq);
rc=srcsrqt("MaryC", "srctest", 0, reqlen, &subreq, &replen,
&statbuf, SRC_NO, &cont);
```

This entry gets long status of the subsystem srctest on the MaryC machine. The subsystem keeps sending status packets until statbuf.srchdr.cont=END.

2. To start a subserver, enter:

```
int cont=NEWREQUEST;
int rc;
short replen;
short reqlen;
struct
{
    struct srchdr srchdr;
    struct statcode statcode[20];
} statbuf;
struct subreq subreq;

subreq.action=START;
subreq.object=1234;
replen=sizeof(statbuf);
reqlen=sizeof(subreq);
rc=srcsrqt("", "", 987, reqlen, &subreq, &replen, &statbuf,
SRC_NO, &cont);
```

This entry starts the subserver with the code point of 1234, but only if the subsystem is already active.

3. To start a subserver and a subsystem, enter:

```
int cont=NEWREQUEST;
int rc;
short replen;
short reqlen;
struct
{
    struct srchdr srchdr;
    struct statcode statcode[20];
} statbuf;
struct subreq subreq;
subreq.action=START;
subreq.object=1234;
replen=sizeof(statbuf);
reqlen=sizeof(subreq);
rc=srcsrqt("", "", 987, reqlen, &subreq, &replen, &statbuf, SRC_YES, &cont);
```

This entry starts the subserver with the code point of 1234. If the subsystem to which this subserver belongs is not active, the subsystem is started.

## Files

<code>/etc/services</code>	Defines sockets and protocols used for Internet services.
<code>/dev/SRC</code>	Specifies the <b>AF_UNIX</b> socket file.
<code>/dev/SRC-unix</code>	Specifies the location for temporary socket files.

## Related Information

The **srcrrqs** (“srcrrqs Subroutine” on page 241) subroutine, **srcsbuf** (“srcsbuf Subroutine” on page 244) subroutine, **srcsrpy** (“srcsrpy Subroutine” on page 250) subroutine, **srcstat** (“srcstat Subroutine” on page 259) subroutine, **srcstathdr** (“srcstathdr Subroutine” on page 264) subroutine, **srcstattxt** (“srcstattxt Subroutine” on page 265) subroutine, **srcstop** (“srcstop Subroutine” on page 266) subroutine, **srcstrt** (“srcstrt Subroutine” on page 269) subroutine.

List of SRC Subroutines, Programming Subsystem Communication with the SRC, System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## srcsrqt\_r Subroutine

### Purpose

Sends a request to a subsystem.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <spc.h>
```

```
srcsrqt_r(Host, SubsystemName, SubsystemPID, RequestLength,  
          SubsystemRequest, ReplyLength, ReplyBuffer, StartItAlso,  
          Continued, SRCHandle)  
char * Host, * SubsystemName;  
char * SubsystemRequest, * ReplyBuffer;  
pid_t SubsystemPID,  
int, StartItAlso, * Continued;  
short RequestLength, * ReplyLength;  
char ** SRCHandle;
```

### Description

The **srcsrqt\_r** subroutine sends a request to a subsystem, waits for a response and returns one or more replies to the caller. The format of the request and the reply is determined by the caller and the subsystem.

**Note:** For each **NEWREQUEST**, the **srcsrqt\_r** subroutine creates its own socket to send a request to the subsystem. The socket that this function opens remains open until an error or an end packet is received.

This system is threadsafe and reentrant.

Two types of continuation are returned by the **srcsrqt\_r** subroutine:

No continuation	<i>ReplyBuffer-&gt;srchdr.continued</i> is set to the <b>END</b> constant.
Reply continuation	<i>ReplyBuffer-&gt;srchdr.continued</i> is not set to the <b>END</b> constant, but to a positive value agreed upon by the calling process and the subsystem. The packet is returned to the caller.

## Parameters

<i>SubsystemPID</i>	The process ID of the subsystem.
<i>Host</i>	Specifies the foreign host on which this subsystem request is to be sent. If the host is null, the request is sent to the subsystem on the local host.
<i>SubsystemName</i>	Specifies the name of the subsystem to which this request is to be sent. You must specify a <i>SubsystemName</i> if you do not specify a <i>SubsystemPID</i> .
<i>RequestLength</i>	Specifies the length, in bytes, of the request to be sent to the subsystem. The maximum length is 2000 bytes.
<i>SubsystemRequest</i>	Points to the subsystem request packet.
<i>ReplyLength</i>	Specifies the maximum length, in bytes, of the reply to be received from the subsystem. On return from the <b>srcsrqt</b> subroutine, the <i>ReplyLength</i> parameter is set to the actual length of the subsystem reply packet.
<i>ReplyBuffer</i>	Points to a buffer for the receipt of the reply packet from the subsystem.
<i>StartItAlso</i>	Specifies whether the subsystem should be started if it is nonactive. When nonzero, the System Resource Controller (SRC) attempts to start a nonactive subsystem, and then passes the request to the subsystem.
<i>Continued</i>	Specifies whether this call to the <b>srcsrqt</b> subroutine is a continuation of a previous request. If the <i>Continued</i> parameter is set to <b>NEWREQUEST</b> , a request for it is sent to the subsystem and the subsystem is notified that a response is expected. Under normal circumstances, the calling process should never set <i>Continued</i> to any value other than <b>NEWREQUEST</b> . The last response from the subsystem will set <i>Continued</i> to <b>END</b> . The caller must continue to call the <b>srcsrqt_r</b> subroutine until <b>END</b> is received. Otherwise, the socket will not be closed and the internal buffers freed. As an alternative, set <i>Continued=SRC_CLOSE</i> to discard the remaining data, close the socket, and free the internal buffers.
<i>SRCHandle</i>	Identifies a request and its associated responses. Set to NULL by the caller for a <b>NEWREQUEST</b> . The <b>srcsrqt_r</b> subroutine saves a value in <i>SRCHandle</i> to allow <b>srcsrqt_r</b> continuation calls to use the same socket and internal buffers. The <i>SRCHandle</i> parameter should not be changed by the caller except for <b>NEWREQUESTs</b> .

## Return Values

If the **srcsrqt\_r** subroutine is successful, the value **SRC\_OK** is returned.

## Error Codes

The **srcsrqt\_r** subroutine fails and returns the corresponding error code if one of the following error conditions is detected:

<b>SRC_BADSOCK</b>	The request could not be passed to the subsystem because of a socket failure.
<b>SRC_CONT</b>	The subsystem uses signals. The request cannot complete.
<b>SRC_DMNA</b>	The SRC daemon is not active.
<b>SRC_INET_AUTHORIZED_HOST</b>	The local host is not in the remote <i>/etc/hosts.equiv</i> file.
<b>SRC_INET_INVALID_HOST</b>	On the remote host, the local host is not known.
<b>SRC_INVALID_USER</b>	The user is not root or group system.
<b>SRC_MMRV</b>	An SRC component could not allocate the memory it needs.

<b>SRC_NOCONTINUE</b>	The <i>Continued</i> parameter was not set to <b>NEWREQUEST</b> , and no continuation is currently active.
<b>SRC_NORPLY</b>	The request timed out waiting for a response.
<b>SRC_NSVR</b>	The subsystem is not active.
<b>SRC_REQLEN2BIG</b>	The <i>RequestLength</i> is greater than the maximum 2000 bytes.
<b>SRC SOCK</b>	There is a problem with SRC socket communications.
<b>SRC_STPG</b>	The request was not passed to the subsystem. The subsystem is stopping.
<b>SRC_UDP</b>	The SRC port is not defined in the <i>/etc/services</i> file.
<b>SRC_UHOST</b>	The foreign host is not known.

## Examples

1. To request long subsystem status, enter:

```
int cont=NEWREQUEST;
int rc;
short replen;
short reqlen;
char *handle;
struct
{
    struct srchdr srchdr;
    struct statcode statcode[20];
} statbuf;
struct subreq subreq;

subreq.action=STATUS;
subreq.object=SUBSYSTEM;
subreq.parm1=LONGSTAT;
strcpy(subreq.objname,"srctest");
replen=sizeof(statbuf);
reqlen=sizeof(subreq);
rc=srcsrqt_r("MaryC", "srctest", 0, reqlen, &subreq, &replen,
&statbuf, SRC_NO, &cont, &handle);
```

This entry gets long status of the subsystem *srctest* on the *MaryC* machine. The subsystem keeps sending status packets until *statbuf.srchdr.cont=END*.

2. To start a subserver, enter:

```
int cont=NEWREQUEST;
int rc;
short replen;
short reqlen;
struct
char *handle;
struct
{
    struct srchdr srchdr;
    struct statcode statcode[20];
} statbuf;
struct subreq subreq;

subreq.action=START;
subreq.object=1234;
replen=sizeof(statbuf);
reqlen=sizeof(subreq);
rc=srcsrqt_r("", "", 987, reqlen, &subreq, &replen, &statbuf,
SRC_NO, &cont, &handle);
```

This entry starts the subserver with the code point of 1234, but only if the subsystem is already active.

3. To start a subserver and a subsystem, enter:

```

int cont=NEWREQUEST;
int rc;
short replen;
short reqlen;
char *handle;
struct
{
    struct srchdr srchdr;
    struct statcode statcode[20];
} statbuf;
struct subreq subreq;
subreq.action=START;
subreq.object=1234;
replen=sizeof(statbuf);
reqlen=sizeof(subreq);
rc=srcsrqt("", "", 987, reqlen, &subreq, &replen, &statbuf, SRC_YES, &cont, &handle);

```

This entry starts the subserver with the code point of 1234. If the subsystem to which this subserver belongs is not active, the subsystem is started.

## Files

<b>/etc/services</b>	Defines sockets and protocols used for Internet services.
<b>/dev/SRC</b>	Specifies the <b>AF_UNIX</b> socket file.
<b>/dev/.SRC-unix</b>	Specifies the location for temporary socket files.

## Related Information

The **src\_err\_msg\_r** (“src\_err\_msg\_r Subroutine” on page 241), **srcsbuf\_r** (“srcsbuf\_r Subroutine” on page 247), **srcrrqs\_r** (“srcrrqs\_r Subroutine” on page 243), **srcstat\_r** (“srcstat\_r Subroutine” on page 262), and **srcstattxt\_r** (“srcstattxt\_r Subroutine” on page 265) subroutines.

List of SRC Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Programming Subsystem Communication with the SRC in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## srcstat Subroutine

### Purpose

Gets short status on one or more subsystems.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <spc.h>
```

```
int srcstat(Host,
SubsystemName, SubsystemPID, ReplyLength, StatusReply, Continued)
char * Host, * SubsystemName;
```

```
int SubsystemPID, * Continued;
short * ReplyLength;
void * StatusReply;
```

## Description

The **srcstat** subroutine sends a short status request to the System Resource Controller (SRC) and returns status for one or more subsystems to the caller.

## Parameters

<i>Host</i>	Specifies the foreign host on which this status action is requested. If the host is null, the status request is sent to the SRC on the local host. The local user must be running as "root". The remote system must be configured to accept remote System Resource Controller requests. That is, the <b>srcmstr</b> daemon (see <b>/etc/inittab</b> ) must be started with the <b>-r</b> flag and the <b>/etc/hosts.equiv</b> or <b>.rhosts</b> file must be configured to allow remote requests.
<i>SubsystemName</i>	Specifies the name of the subsystem on which to get short status. To get status of all subsystems, use the <b>SRCALLSUBSYS</b> constant. To get status of a group of subsystems, the <i>SubsystemName</i> parameter must start with the <b>SRCGROUP</b> constant, followed by the name of the group for which you want status appended. If you specify a null <i>SubsystemName</i> parameter, you must specify a <i>SubsystemPID</i> parameter.
<i>SubsystemPID</i>	Specifies the PID of the subsystem on which to get status as returned by the <b>srcstat</b> subroutine. You must specify the <i>SubsystemPID</i> parameter if multiple instances of the subsystem are active and you request a long subsystem status or subserver status. If you specify a null <i>SubsystemPID</i> parameter, you must specify a <i>SubsystemName</i> parameter.
<i>ReplyLength</i>	Specifies size of a <b>srchdr</b> structure plus the number of <b>statcode</b> structures times the size of one <b>statcode</b> structure. On return from the <b>srcstat</b> subroutine, this value is updated.
<i>StatusReply</i>	Specifies a pointer to a structure containing first element as struct <b>srchdr</b> and secondary element as struct <b>statcode</b> (both defined in <b>spc.h</b> file) array that receives the status reply for the requested subsystem. The first element of the returned <b>statcode</b> array contains the status title line. The number of <b>statcode</b> structures array items depends on the number of subsystems user queried.
<i>Continued</i>	Specifies whether this call to the <b>srcstat</b> subroutine is a continuation of a previous status request. If the <i>Continued</i> parameter is set to <b>NEWREQUEST</b> , a request for short subsystem status is sent to the SRC and <b>srcstat</b> waits for the first status response. The calling process should never set <i>Continued</i> to a value other than <b>NEWREQUEST</b> . The last response for the SRC sets <i>Continued</i> to <b>END</b> .

## Return Values

If the **srcstat** subroutine succeeds, it returns a value of 0. An error code is returned if the subroutine is unsuccessful.

## Error Codes

The **srcstat** subroutine fails if one or more of the following are true:

<b>SRC_DMNA</b>	The SRC daemon is not active.
<b>SRC_INET_AUTHORIZED_HOST</b>	The local host is not in the remote <b>/etc/hosts.equiv</b> file.
<b>SRC_INET_INVALID_HOST</b>	On the remote host, the local host is not known.
<b>SRC_INVALID_USER</b>	The user is not root or group system.
<b>SRC_MMRV</b>	An SRC component could not allocate the memory it needs.
<b>SRC_NOCONTINUE</b>	<i>Continued</i> was not set to <b>NEWREQUEST</b> and no continuation is currently active.
<b>SRC_NORPLY</b>	The request timed out waiting for a response.
<b>SRC_SOCKET</b>	There is a problem with SRC socket communications.
<b>SRC_UDP</b>	The SRC port is not defined in the <b>/etc/services</b> file.

## Examples

1. To request the status of a subsystem, enter:

```
intcont=NEWREQUEST;
struct {
    struct srchdr srchdr
    struct statcode statcode[6];
} status;
short replen=sizeof(status);

srcstat("MaryC", "srctest", 0, &replen, &status, &cont);
```

This entry requests short status of all instances of the subsystem `srctest` on the `MaryC` machine.

2. To request the status of all subsystems, enter:

```
int cont=NEWREQUEST;
struct {
    struct srchdr srchdr;
    struct statcode statcode[80];
} status;
short replen=sizeof(status);

srcstat("", SRCALLSUBSYS, 0, &replen, &status, &cont);
```

This entry requests short status of all subsystems on the local machine.

3. To request the status for a group of subsystems, enter:

```
int cont=NEWREQUEST;
struct struct {
    struct srchdr srchdr;
    struct statcode statcode[30];
} status;
short replen=sizeof(status), rep_num;
char subsysname[30];

strcpy(subsysname, SRCGROUP);
strcat(subsysname, "tcpip");
srcstat("", subsysname, 0, &replen, &status, &cont);

rep_num = (replen - sizeof(struct srchdr)) / sizeof(struct statcode);

for (i = 0; i < rep_num; i++)
    printf("objtype %d status %d objname %s objtext %s\n",
        status.statcode[i].objtype, status.statcode[i].status,
        status.statcode[i].objname, status.statcode[i].objtext);
```

This entry requests short status of all members of the subsystem group `tcpip` on the local machine, and displays the query results on **stdout**.

## Files

<code>/etc/services</code>	Defines the sockets and protocols used for Internet services.
<code>/dev/SRC</code>	Specifies the <b>AF_UNIX</b> socket file.
<code>/dev/.SRC-unix</code>	Specifies the location for temporary socket files.

## Related Information

The **srcrrqs** (“`srcrrqs` Subroutine” on page 241) subroutine, **srcsbuf** (“`srcsbuf` Subroutine” on page 244) subroutine, **srcsrpy** (“`srcsrpy` Subroutine” on page 250) subroutine, **srcsrqt** (“`srcsrqt` Subroutine” on page 253)

page 253) subroutine, **srcstathdr** (“srcstathdr Subroutine” on page 264) subroutine, **srcstattxt** (“srcstattxt Subroutine” on page 265) subroutine, **srcstop** (“srcstop Subroutine” on page 266) subroutine, **srcstr** (“srcstr Subroutine” on page 269) subroutine.

List of SRC Subroutines, Programming Subsystem Communication with the SRC, System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## srcstat\_r Subroutine

### Purpose

Gets short status on a subsystem.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <spc.h>
```

```
int srcstat_r(Host, SubsystemName, SubsystemPID, ReplyLength,
              StatusReply, Continued, SRCHandle)
char * Host, * SubsystemName;
pid_t SubsystemPID;
int * Continued;
short * ReplyLength;
struct statrep * StatusReply;
char ** SRCHandle;
```

### Description

The **srcstat\_r** subroutine sends a short status request to the System Resource Controller (SRC) and returns status for one or more subsystems to the caller. This subroutine is threadsafe and reentrant.

### Parameters

<i>Host</i>	Specifies the foreign host on which this status action is requested. If the host is null, the status request is sent to the SRC on the local host.
<i>SubsystemName</i>	Specifies the name of the subsystem on which to get short status. To get status of all subsystems, use the <b>SRCALLSUBSYS</b> constant. To get status of a group of subsystems, the <i>SubsystemName</i> parameter must start with the <b>SRCGROUP</b> constant, followed by the name of the group for which you want status appended. If you specify a null <i>SubsystemName</i> parameter, you must specify a <i>SubsystemPID</i> parameter.
<i>SubsystemPID</i>	Specifies the PID of the subsystem on which to get status as returned by the <b>srcstat_r</b> subroutine. You must specify the <i>SubsystemPID</i> parameter if multiple instances of the subsystem are active and you request a long subsystem status or subserver status. If you specify a null <i>SubsystemPID</i> parameter, you must specify a <i>SubsystemName</i> parameter.
<i>ReplyLength</i>	Specifies size of a <b>srchdr</b> structure plus the number of <b>statcode</b> structures times the size of one <b>statcode</b> structure. On return from the <b>srcstat_r</b> subroutine, this value is updated.
<i>StatusReply</i>	Specifies a pointer to a <b>statrep</b> code structure containing a <b>statcode</b> array that receives the status reply for the requested subsystem. The first element of the returned <b>statcode</b> array contains the status title line. The <b>statcode</b> structure is defined in the <b>spc.h</b> file.

*Continued*

Specifies whether this call to the **srcstat\_r** subroutine is a continuation of a previous status request. If the *Continued* parameter is set to **NEWREQUEST**, a request for short subsystem status is sent to the SRC and **srcstat\_r** waits for the first status response. During **NEWREQUEST** processing, the **srcstat\_r** subroutine opens a socket, allocates internal buffers, and saves a value in *SRCHandle*. In normal circumstances, the calling process should never set *Continued* to a value other than **NEWREQUEST**. When the **srcstat\_r** subroutine returns with *Continued=STATCONTINUED*, call **srcstat\_r** without changing the *Continued* and *SRCHandle* parameters to receive additional data. The last response from the SRC sets *Continued* to **END**. The caller must continue to call **srcstat\_r** until **END** is received. Otherwise, the socket will not be closed and the internal buffers freed. As an alternative, call **srcstat\_r** with *Continued=STATCONTINUED* to discard the remaining data, close the socket, and free the internal buffers.

*SRCHandle*

Identifies a request and its associated responses. Set to NULL by the caller for a **NEWREQUEST**. The **srcstat\_r** subroutine saves a value in *SRCHandle* to allow subsequent **srcstat\_r** calls to use the same socket and internal buffers. The *SRCHandle* parameter should not be changed by the caller except for **NEWREQUESTS**.

## Return Values

If the **srcstat\_r** subroutine succeeds, it returns a value of 0. An error code is returned if the subroutine is unsuccessful.

## Error Codes

The **srcstat\_r** subroutine fails and returns the corresponding error code if one of the following error conditions is detected:

<b>SRC_DMNA</b>	The SRC daemon is not active.
<b>SRC_INET_AUTHORIZED_HOST</b>	The local host is not in the remote <i>/etc/hosts.equiv</i> file.
<b>SRC_INET_INVALID_HOST</b>	On the remote host, the local host is not known.
<b>SRC_INVALID_USER</b>	The user is not root or group system.
<b>SRC_MMRV</b>	An SRC component could not allocate the memory it needs.
<b>SRC_NOCONTINUE</b>	<i>Continued</i> was not set to <b>NEWREQUEST</b> and no continuation is currently active.
<b>SRC_NORPLY</b>	The request timed out waiting for a response.
<b>SRC_SOCKET</b>	There is a problem with SRC socket communications.
<b>SRC_UDP</b>	The SRC port is not defined in the <i>/etc/services</i> file.
<b>SRC_UHOST</b>	The foreign host is not known.

## Examples

1. To request the status of a subsystem, enter:

```
int cont=NEWREQUEST;
struct statcode statcode[20];
short replen=sizeof(statcode);
char *handle;

srcstat_r("MaryC","srctest",0,&replen,statcode, &cont, &handle);
```

This entry requests short status of all instances of the subsystem *srctest* on the *MaryC* machine.

2. To request the status of all subsystems, enter:

```
int cont=NEWREQUEST;
struct statcode statcode[20];
short replen=sizeof(statcode);
char *handle;

srcstat_r("",SRCALLSUBSYS,0,&replen,statcode, &cont, &handle);
```

This entry requests short status of all subsystems on the local machine.

3. To request the status for a group of subsystems, enter:

```
int cont=NEWREQUEST;
struct statcode statcode[20];
short replen=sizeof(statcode);
char subsysname[30];
char *handle;

strcpy(subsysname,SRCGROUP);
strcat(subsysname,"tcpip");
srcstat_r("",subsysname,0,&replen,statcode, &cont, &handle);
```

This entry requests short status of all members of the subsystem group tcpip on the local machine.

## Files

<code>/etc/services</code>	Defines the sockets and protocols used for Internet services.
<code>/dev/SRC</code>	Specifies the <b>AF_UNIX</b> socket file.
<code>/dev.SRC-unix</code>	Specifies the location for temporary socket files.

## Related Information

The **src\_err\_msg\_r** (“src\_err\_msg\_r Subroutine” on page 241), **srcsbuf\_r** (“srcsbuf\_r Subroutine” on page 247), **srcsrqt\_r** (“srcsrqt\_r Subroutine” on page 256), **srcrrqs\_r** (“srcrrqs\_r Subroutine” on page 243), and **srcstattxt\_r** (“srcstattxt\_r Subroutine” on page 265) subroutines.

---

## srcstathdr Subroutine

### Purpose

Gets the title line of the System Resource Controller (SRC) status text.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
void srcstathdr ( Title1, Title2)
char *Title1, *Title2;
```

### Description

The **srcstathdr** subroutine retrieves the title line, or header, of the SRC status text.

### Parameters

<i>Title1</i>	Specifies the objname field of a <b>statcode</b> structure. The subsystem name title is placed here.
<i>Title2</i>	Specifies the objtext field of a <b>statcode</b> structure. The remaining titles are placed here.

### Return Values

The subsystem name title is returned in the *Title1* parameter. The remaining titles are returned in the *Title2* parameter.

## Related Information

The **srcrrqs** (“srcrrqs Subroutine” on page 241) subroutine, **srcsbuf** (“srcsbuf Subroutine” on page 244) subroutine, **srcsrpy** (“srcsrpy Subroutine” on page 250) subroutine, **srcsrqt** (“srcsrqt Subroutine” on page 253) subroutine, **srcstat** (“srcstat Subroutine” on page 259) subroutine, **srcstattxt** (“srcstattxt Subroutine”) subroutine, **srcstop** (“srcstop Subroutine” on page 266) subroutine, **srcstrt** (“srcstrt Subroutine” on page 269) subroutine.

List of SRC Subroutines, Programming Subsystem Communication with the SRC, System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## srcstattxt Subroutine

### Purpose

Gets the System Resource Controller (SRC) status text representation for a status code.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
char *srcstattxt ( StatusCode)  
short StatusCode;
```

### Description

The **srcstattxt** subroutine, given an SRC status code, gets the text representation and returns a pointer to this text.

### Parameters

*StatusCode*            Specifies an SRC status code to be translated into meaningful text.

### Return Values

The **srcstattxt** subroutine returns a pointer to the text representation of a status code.

## Related Information

The **srcrrqs** (“srcrrqs Subroutine” on page 241) subroutine, **srcsbuf** (“srcsbuf Subroutine” on page 244) subroutine, **srcsrpy** (“srcsrpy Subroutine” on page 250) subroutine, **srcsrqt** (“srcsrqt Subroutine” on page 253) subroutine, **srcstat** (“srcstat Subroutine” on page 259) subroutine, **srcstathdr** (“srcstathdr Subroutine” on page 264) subroutine, **srcstop** (“srcstop Subroutine” on page 266) subroutine, **srcstrt** (“srcstrt Subroutine” on page 269) subroutine.

List of SRC Subroutines, Programming Subsystem Communication with the SRC, System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## srcstattxt\_r Subroutine

### Purpose

Gets the status text representation for an SRC status code.

## Library

System Resource Controller Library (**libsrc.a**)

## Syntax

```
#include <spc.h>
```

```
char *srcstattxt_r (StatusCode, Text)  
short StatusCode;  
char *Text;
```

## Description

The **srcstattxt\_r** subroutine, given an SRC status code, gets the text representation and returns it in a caller-supplied buffer. This routine is threadsafe and reentrant.

## Parameters

<i>StatusCode</i>	Specifies an SRC status code to be translated into meaningful text.
<i>Text</i>	Points to a caller-supplied buffer where the text will be returned. If the length of the text is unknown, the maximum text length can be used when allocating the buffer. The maximum text length is <b>SRC_STAT_MAX</b> in <b>/usr/include/spc.h</b> (64 bytes).

## Return Values

Upon successful completion, the **srcstattxt\_r** subroutine returns the address of the caller-supplied buffer. Otherwise, no text is returned and the subroutine returns NULL.

## Related Information

The **src\_err\_msg\_r** (“src\_err\_msg\_r Subroutine” on page 241), **srcsbuf\_r** (“srcsbuf\_r Subroutine” on page 247), **srcsrqt\_r** (“srcsrqt\_r Subroutine” on page 256), **srcrrqs\_r** (“srcrrqs\_r Subroutine” on page 243), and **srcstat\_r** (“srcstat\_r Subroutine” on page 262) subroutines.

List of SRC Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Programming Subsystem Communication with the SRC in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## srcstop Subroutine

### Purpose

Stops a System Resource Controller (SRC) subsystem.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include <spc.h>
```

```

srcstop(Host, SubsystemName, SubsystemPID, StopType)
srcstop(ReplyLength, ServerReply, StopFrom)
char * Host, * SubsystemName;
int SubsystemPID, StopFrom;
short StopType, * ReplyLength;
struct srcrep * ServerReply;

```

## Description

The **srcstop** subroutine sends a stop subsystem request to a subsystem and waits for a stop reply from the System Resource Controller (SRC) or the subsystem. The **srcstop** subroutine can only stop a subsystem that was started by the SRC.

## Parameters

<i>Host</i>	Specifies the foreign host on which this stop action is requested. If the host is the null value, the request is sent to the SRC on the local host. The local user must be running as "root". The remote system must be configured to accept remote System Resource Controller requests. That is, the <b>srcmstr</b> daemon (see <b>/etc/inittab</b> ) must be started with the <b>-r</b> flag and the <b>/etc/hosts.equiv</b> or <b>.rhosts</b> file must be configured to allow remote requests.
<i>SubsystemName</i>	Specifies the name of the subsystem to stop.
<i>SubsystemPID</i>	Specifies the process ID of the system to stop as returned by the <b>srcstr</b> subroutine. If you specify a null <i>SubsystemPID</i> parameter, you must specify a <i>SubsystemName</i> parameter.
<i>StopType</i>	Specifies the type of stop requested of the subsystem. If this parameter is null, a normal stop is assumed. The <i>StopType</i> parameter must be one of the following values:  <b>CANCEL</b> Requires a quick stop of the subsystem. The subsystem is sent a <b>SIGTERM</b> signal. After the wait time defined in the subsystem object, the SRC issues a <b>SIGKILL</b> signal to the subsystem. This waiting period allows the subsystem to clean up all its resources and terminate. The stop reply is returned by the SRC.  <b>FORCE</b> Requests a quick stop of the subsystem and all its subservers. The stop reply is returned by the SRC for subsystems that use signals and by the subsystem for other communication types.  <b>NORMAL</b> Requests the subsystem to terminate after all current subsystem activity has completed. The stop reply is returned by the SRC for subsystems that use signals and by the subsystem for other communication types.
<i>ReplyLength</i>	Specifies the maximum length, in bytes, of the stop reply. On return from the <b>srcstop</b> subroutine, this field is set to the actual length of the subsystem reply packet received.
<i>ServerReply</i>	Points to an <b>svrreply</b> structure that will receive the subsystem stop reply.
<i>StopFrom</i>	Specifies whether the <b>srcstop</b> subroutine is to display stop results to standard output. If the <i>StopFrom</i> parameter is set to <b>SSHELL</b> , the stop results are displayed to standard output and the <b>srcstop</b> subroutine returns successfully. If the <i>StopFrom</i> parameter is set to <b>SDAEMON</b> , the stop results are not displayed to standard output, but are passed back to the caller.

## Return Values

Upon successful completion, the **srcstop** subroutine returns **SRC\_OK** or **SRC\_STPOK**.

## Error Codes

The **srcstop** subroutine fails if one or more of the following are true:

**SRC\_BADFSIG** The stop force signal is an invalid signal.

<b>SRC_BADNSIG</b>	The stop normal signal is an invalid signal.
<b>SRC_BADSOCK</b>	The stop request could not be passed to the subsystem on its communication socket.
<b>SRC_DMNA</b>	The SRC daemon is not active.
<b>SRC_INET_AUTHORIZED_HOST</b>	The local host is not in the remote <b>/etc/hosts.equiv</b> file.
<b>SRC_INET_INVALID_HOST</b>	On the remote host, the local host is not known.
<b>SRC_INVALID_USER</b>	The user is not root or group system.
<b>SRC_MMRV</b>	An SRC component could not allocate the memory it needs.
<b>SRC_NORPLY</b>	The request timed out waiting for a response.
<b>SRC_NOTROOT</b>	The SRC daemon is not running as root.
<b>SRC SOCK</b>	There is a problem with SRC socket communications.
<b>SRC_STPG</b>	The request was not passed to the subsystem. The subsystem is stopping.
<b>SRC_SVND</b>	The subsystem is unknown to the SRC daemon.
<b>SRC_UDP</b>	The remote SRC port is not defined in the <b>/etc/services</b> file.
<b>SRC_UHOST</b>	The foreign host is not known.
<b>SRC_PARM</b>	Invalid parameter passed.

## Examples

1. To stop all instances of a subsystem, enter:

```
int rc;
struct svrreply svrreply;
short replen=sizeof(svrreply);

rc=srcstop("MaryC","srctest",0,FORCE,&replen,&svrreply,SDAEMON);
```

This request stops a subsystem with a stop type of FORCE for all instances of the subsystem srctest on the MaryC machine and does not print a message to standard output about the status of the stop.

2. To stop a single instance of a subsystem, enter:

```
struct svrreply svrreply;
short replen=sizeof(svrreply);

rc=srcstop("", "", 999,CANCEL,&replen,&svrreply,SSHLL);
```

This request stops a subsystem with a stop type of CANCEL, with the process ID of 999 on the local machine and prints a message to standard output about the status of the stop.

## Files

<b>/etc/services</b>	Defines sockets and protocols used for Internet services.
<b>/dev/SRC</b>	Specifies the <b>AF_UNIX</b> socket file.
<b>/dev/.SRC-unix</b>	Specifies the location for temporary socket files.

## Related Information

The **srcrrqs** (“srcrrqs Subroutine” on page 241) subroutine, **srcsbuf** (“srcsbuf Subroutine” on page 244) subroutine, **srcsrpy** (“srcsrpy Subroutine” on page 250) subroutine, **srcsrqt** (“srcsrqt Subroutine” on page 253) subroutine, **srcstat** (“srcstat Subroutine” on page 259) subroutine, **srcstathdr** (“srcstathdr Subroutine” on page 264) subroutine, **srcstattxt** (“srcstattxt Subroutine” on page 265) subroutine, **srcstrt** (“srcstrt Subroutine” on page 269) subroutine.

List of SRC Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Programming Subsystem Communication with the SRC in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## srcstrt Subroutine

### Purpose

Starts a System Resource Controller (SRC) subsystem.

### Library

System Resource Controller Library (**libsrc.a**)

### Syntax

```
#include<spc.h>
```

```
srcstrt (Host, SubsystemName, Environment, Arguments, Restart, StartFrom)
```

```
char * Host, * SubsystemName;
```

```
char * Environment, * Arguments;
```

```
unsigned int Restart;
```

```
int StartFrom;
```

### Description

The **srcstrt** subroutine sends a start subsystem request packet and waits for a reply from the System Resource Controller (SRC).

### Parameters

<i>Host</i>	Specifies the foreign host on which this start subsystem action is requested. If the host is null, the request is sent to the SRC on the local host. The local user must be running as "root". The remote system must be configured to accept remote System Resource Controller requests. That is, the <b>srcmstr</b> daemon (see <b>/etc/inittab</b> ) must be started with the <b>-r</b> flag and the <b>/etc/hosts.equiv</b> or <b>.rhosts</b> file must be configured to allow remote requests.
<i>SubsystemName</i>	Specifies the name of the subsystem to start.
<i>Environment</i>	Specifies a string that is placed in the subsystem environment when the subsystem is executed. The combined values of the <i>Environment</i> and <i>Arguments</i> parameters cannot exceed a maximum of 2400 characters. Otherwise, the <b>srcstrt</b> subroutine will fail. The environment string is parsed by the SRC according to the same rules used by the shell. For example, quoted strings are passed as a single <i>Environment</i> value, and blanks outside a quoted string delimit each environment value.
<i>Arguments</i>	Specifies a string that is passed to the subsystem when the subsystem is executed. The string is parsed from the command line and appended to the command line arguments from the subsystem object class. The combined values of the <i>Environment</i> and <i>Arguments</i> parameters cannot exceed a maximum of 2400 characters. Otherwise, the <b>srcstrt</b> subroutine will fail. The command argument is parsed by the SRC according to the same rules used by the shell. For example, quoted strings are passed as a single argument, and blanks outside a quoted string delimit each argument.

<i>Restart</i>	Specifies override on subsystem restart. If the <i>Restart</i> parameter is set to <b>SRCNO</b> , the subsystem's restart definition from the subsystem object class is used. If the <i>Restart</i> parameter is set to <b>SRCYES</b> , the restart of a subsystem is not attempted if it terminates abnormally.
<i>StartFrom</i>	Specifies whether the <b>srcstrt</b> subroutine is to display start results to standard output. If the <i>StartFrom</i> parameter is set to <b>SSHELL</b> , the start results are displayed to standard output, and the <b>srcstrt</b> subroutine always returns successfully. If the <i>StartFrom</i> parameter is set to <b>SDAEMON</b> , the start results are not displayed to standard output but are passed back to the caller.

## Return Values

When the *StartFrom* parameter is set to **SSHELL**, the **srcstrt** subroutine returns the value **SRC\_OK**. Otherwise, it returns the subsystem process ID.

## Error Codes

The **srcstrt** subroutine fails if any of the following are true:

<b>SRC_AUDITID</b>	The audit user ID is invalid.
<b>SRC_DMNA</b>	The SRC daemon is not active.
<b>SRC_FEXE</b>	The subsystem could not be <b>forked</b> and <b>execed</b> .
<b>SRC_INET_AUTHORIZED_HOST</b>	The local host is not in the remote <b>/etc/hosts.equiv</b> file.
<b>SRC_INET_INVALID_HOST</b>	On the remote host, the local host is not known.
<b>SRC_INVALID_USER</b>	The user is not root or group system.
<b>SRC_INPT</b>	The subsystem standard input file could not be established.
<b>SRC_MMRV</b>	An SRC component could not allocate the memory it needs.
<b>SRC_MSGQ</b>	The subsystem message queue could not be created.
<b>SRC_MULT</b>	Multiple instance of the subsystem are not allowed.
<b>SRC_NORPLY</b>	The request timed out waiting for a response.
<b>SRC_OUT</b>	The subsystem standard output file could not be established.
<b>SRC_PIPE</b>	A pipe could not be established for the subsystem.
<b>SRC_SERR</b>	The subsystem standard error file could not be established.
<b>SRC_SUBSOCK</b>	The subsystem communication socket could not be created.
<b>SRC_SUBSYSID</b>	The system user ID is invalid.
<b>SRC_SOCK</b>	There is a problem with SRC socket communications.
<b>SRC_SVND</b>	The subsystem is unknown to the SRC daemon.
<b>SRC_UDP</b>	The SRC port is not defined in the <b>/etc/services</b> header file.
<b>SRC_UHOST</b>	The foreign host is not known.

## Examples

1. To start a subsystem passing the *Environment* and *Arguments* parameters, enter:

```
rc=srcstrt("", "srctest", "HOME=/tmpTERM=ibm6155",
"-z\"thezflagargument\"", SRC_YES, SSHELL);
```

This starts the **srctest** subsystem on the local host, placing **HOME=/tmp**, **TERM=ibm6155** in the environment and using **-z** and **thezflagargument** as two arguments to the subsystem. This also displays the results of the start command to standard output and allows the SRC to restart the subsystem should it end abnormally.

2. To start a subsystem on a foreign host, enter:

```
rc=srcstrt("MaryC", "srctest", "", "", SRC_NO, SDAEMON);
```

This starts the **srctest** subsystem on the **MaryC** machine. This does not display the results of the start command to standard output and does not allow the SRC to restart the subsystem should it end abnormally.

## Files

<code>/etc/services</code>	Defines sockets and protocols used for Internet services.
<code>/dev/SRC</code>	Specifies the <b>AF_UNIX</b> socket file.
<code>/dev.SRC-unix</code>	Specifies the location for temporary socket files.

## Related Information

The **srcrrqs** (“srcrrqs Subroutine” on page 241) subroutine, **srcsbuf** (“srcsbuf Subroutine” on page 244) subroutine, **srcsrpy** (“srcsrpy Subroutine” on page 250) subroutine, **srcsrqt** (“srcsrqt Subroutine” on page 253) subroutine, **srcstat** (“srcstat Subroutine” on page 259) subroutine, **srcstathdr** (“srcstathdr Subroutine” on page 264) subroutine, **srcstattxt** (“srcstattxt Subroutine” on page 265) subroutine, **srcstop** (“srcstop Subroutine” on page 266) subroutine.

List of SRC Subroutines, Programming Subsystem Communication with the SRC, System Resource Controller (SRC) Overview for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## ssignal or gsignal Subroutine

### Purpose

Implements a software signal facility.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
void (*ssignal ( Signal, Action))( )  
int Signal;  
void (*Action)( );  
int gsignal (Signal)  
int Signal;
```

### Description

**Attention:** Do not use the **ssignal** or **gsignal** subroutine in a multithreaded environment.

The **ssignal** and **gsignal** subroutines implement a software facility similar to that of the **signal** and **kill** subroutines. However, there is no connection between the two facilities. User programs can use the **ssignal** and **gsignal** subroutines to handle exceptional processing within an application. The **signal** subroutine and related subroutines handle system-defined exceptions.

The software signals available are associated with integers in the range 1 through 16. Other values are reserved for use by the C library and should not be used.

The **ssignal** subroutine associates the procedure specified by the *Action* parameter with the software signal specified by the *Signal* parameter. The **gsignal** subroutine raises the *Signal*, causing the procedure specified by the *Action* parameter to be taken.

The *Action* parameter is either a pointer to a user-defined subroutine, or one of the constants **SIG\_DFL** (default action) and **SIG\_IGN** (ignore signal). The **ssignal** subroutine returns the procedure that was

previously established for that signal. If no procedure was established before, or if the signal number is illegal, then the **signal** subroutine returns the value of **SIG\_DFL**.

The **gsignal** subroutine raises the signal specified by the *Signal* parameter by doing the following:

- If the procedure for the *Signal* parameter is **SIG\_DFL**, the **gsignal** subroutine returns a value of 0 and takes no other action.
- If the procedure for the *Signal* parameter is **SIG\_IGN**, the **gsignal** subroutine returns a value of 1 and takes no other action.
- If the procedure for the *Signal* parameter is a subroutine, the *Action* value is reset to the **SIG\_DFL** procedure and the subroutine is called, with the *Signal* value passed as its parameter. The **gsignal** subroutine returns the value returned by the signal-handling routine.
- If the *Signal* parameter specifies an illegal value or if no procedure is specified for that signal, the **gsignal** subroutine returns a value of 0 and takes no other action.

## Parameters

*Signal*        Specifies a signal.  
*Action*       Specifies a procedure.

## Related Information

The **kill** or **killpg** subroutine, **signal** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine.

---

## statacl or fstatacl Subroutine

### Purpose

Retrieves the access control information for a file.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int statacl (Path, Command, ACL, ACLSize)  
char * Path;  
int Command;  
struct acl * ACL;  
int ACLSize;
```

```
int fstatacl (FileDescriptor, Command, ACL, ACLSize)  
int FileDescriptor;  
int Command;  
struct acl * ACL;  
int ACLSize;
```

### Description

The **statacl** and **fstatacl** subroutines return the access control information for a file system object.

## Parameters

<i>Path</i>	Specifies a pointer to the path name of a file.
<i>FileDescriptor</i>	Specifies the file descriptor of an open file.
<i>Command</i>	Specifies the mode of the path interpretation for <i>Path</i> , specifically whether to retrieve information about a symbolic link or mount point. Valid values for the <i>Command</i> parameter are defined in the <b>stat.h</b> file and include: <ul style="list-style-type: none"><li>• <b>STX_LINK</b></li><li>• <b>STX_MOUNT</b></li><li>• <b>STX_NORMAL</b></li></ul>
<i>ACL</i>	Specifies a pointer to a buffer to contain the Access Control List (ACL) of the file system object. The format of an ACL is defined in the <b>sys/acl.h</b> file and includes the following members: <b>acl_len</b> Size of the Access Control List (ACL). <b>Note:</b> The entire ACL for a file cannot exceed one memory page (4096 bytes). <b>acl_mode</b> File mode. <b>Note:</b> The valid values for the <b>acl_mode</b> are defined in the <b>sys/mode.h</b> file. <b>u_access</b> Access permissions for the file owner. <b>g_access</b> Access permissions for the file group. <b>o_access</b> Access permissions for default class <i>others</i> . <b>acl_ext[ ]</b> An array of the extended entries for this access control list.  The members for the base ACL (owner, group, and others) may contain the following bits, which are defined in the <b>sys/access.h</b> file: <b>R_ACC</b> Allows read permission. <b>W_ACC</b> Allows write permission. <b>X_ACC</b> Allows execute or search permission.
<i>ACLSize</i>	Specifies the size of the buffer to contain the ACL. If this value is too small, the first word of the ACL is set to the size of the buffer needed.

## Return Values

On successful completion, the **statacl** and **fstatacl** subroutines return a value of 0. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **statacl** subroutine fails if one or more of the following are true:

<b>ENOTDIR</b>	A component of the <i>Path</i> prefix is not a directory.
<b>ENOENT</b>	A component of the <i>Path</i> does not exist or has the <b>disallow truncation</b> attribute (see the <b>ulimit</b> subroutine).
<b>ENOENT</b>	The <i>Path</i> parameter was null.
<b>EACCESS</b>	Search permission is denied on a component of the <i>Path</i> prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points to a location outside of the allocated address space of the process.

<b>ESTALE</b>	The process' root or current directory is located in a virtual file system that has been unmounted.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENAMETOOLONG</b>	A component of the <i>Path</i> parameter exceeded 255 characters, or the entire <i>Path</i> parameter exceeded 1023 characters.

The **fstatacl** subroutine fails if the following is true:

**EBADF** The file descriptor *FileDescriptor* is not valid.

The **statacl** or **fstatacl** subroutine fails if one or more of the following are true:

<b>EFAULT</b>	The <i>ACL</i> parameter points to a location outside of the allocated address space of the process.
<b>EINVAL</b>	The <i>Command</i> parameter is not a value of <b>STX_LINK</b> , <b>STX_MOUNT</b> , <b>STX_NORMAL</b> .
<b>ENOSPC</b>	The <i>ACLSize</i> parameter indicates the buffer at <i>ACL</i> is too small to hold the Access Control List. In this case, the first word of the buffer is set to the size of the buffer required.
<b>EIO</b>	An I/O error occurred during the operation.

If Network File System (NFS) is installed on your system, the **statacl** and **fstatacl** subroutines can also fail if the following is true:

**ETIMEDOUT** The connection timed out.

## Related Information

The **chacl** subroutine, **stat** (“statx, stat, lstat, fstatx, fstat, fullstat, ffullstat, stat64, lstat64, or fstat64 Subroutine” on page 277) subroutine.

The **acl\_chg** subroutine, **acl\_get** subroutine, **acl\_put** subroutine, **acl\_set** subroutine.

The **aclget** command, **aclput** command, **chmod** command.

List of Security and Auditing Subroutines and Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## statfs, fstatfs, or ustat Subroutine

### Purpose

Gets file system statistics.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```

int statfs ( Path, StatusBuffer)
char *Path;
struct statfs *StatusBuffer;

int fstatfs ( FileDescriptor, StatusBuffer)
int FileDescriptor;
struct statfs *StatusBuffer;
#include <sys/types.h>
#include <ustat.h>

int ustat ( Device, Buffer)
dev_t Device;
struct ustat *Buffer;

```

## Description

The **statfs** and **fstatfs** subroutines return information about the mounted file system that contains the file named by the *Path* or *FileDescriptor* parameters. The returned information is in the format of a **statfs** structure, described in the **sys/statfs.h** file.

The **ustat** subroutine also returns information about a mounted file system identified by *Device*. This device identifier is for any given file and can be determined by examining the *st\_dev* field of the **stat** structure defined in the **sys/stat.h** file. The returned information is in the format of a **ustat** structure, described in the **ustat.h** file. The **ustat** subroutine is superseded by the **statfs** and **fstatfs** subroutines. Use one of these (**statfs** and **fstatfs**) subroutines instead.

## Parameters

<i>Path</i>	The path name of any file within the mounted file system.
<i>FileDescriptor</i>	A file descriptor obtained by a successful <b>open</b> or <b>fcntl</b> subroutine. A file descriptor is a small positive integer used instead of a file name.
<i>StatusBuffer</i>	A pointer to a <b>statfs</b> buffer for the returned information from the <b>statfs</b> or <b>fstatfs</b> subroutine.
<i>Device</i>	The ID of the device. It corresponds to the <i>st_rdev</i> field of the structure returned by the <b>stat</b> subroutine. The <b>stat</b> subroutine and the <b>sys/stat.h</b> file provide more information about the device driver.
<i>Buffer</i>	A pointer to a <b>ustat</b> buffer to hold the returned information.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **statfs**, **fstatfs**, and **ustat** subroutines fail if the following is true:

**EFAULT** The *Buffer* parameter points to a location outside of the allocated address space of the process.

The **fstatfs** subroutine fails if the following is true:

**EBADF** The *FileDescriptor* parameter is not a valid file descriptor.

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

The **statfs** subroutine can be unsuccessful for other reasons. For a list of additional errors, see "Base Operating System Error Codes For Services That Require Path-Name Resolution".



<b>EBADF</b>	The file referred to by the <i>FileDescriptor</i> parameter is not an open file descriptor.
<b>EIO</b>	An I/O error occurred while reading from the filesystem.
<b>ELOOP</b>	Too many symbolic links encountered in translating path.
<b>ENAMETOOLONG</b>	The length of the pathname exceeds <b>PATH_MAX</b> , or name component is longer than <b>NAME_MAX</b> .
<b>ENOENT</b>	The file referred to by the <i>Path</i> parameter does not exist.
<b>ENOMEM</b>	A memory allocation failed during information retrieval.
<b>ENOTDIR</b>	A component of the <i>Path</i> parameter prefix is not a directory.
<b>EOVERFLOW</b>	One of the values to be returned cannot be represented correctly in the structure pointed to by <b>buf</b> .

## Related Information

The **stat** (“statx, stat, lstat, fstatx, fstat, fullstat, ffullstat, stat64, lstat64, or fstat64 Subroutine”) subroutine, **statfs** (“statfs, fstatfs, or ustat Subroutine” on page 274) subroutine.

---

## statx, stat, lstat, fstatx, fstat, fullstat, ffullstat, stat64, lstat64, or fstat64 Subroutine

### Purpose

Provides information about a file.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int stat ( Path, Buffer)
```

```
const char *Path;
```

```
struct stat *Buffer;
```

```
int lstat (Path, Buffer)
```

```
const char *Path;
```

```
struct stat *Buffer;
```

```
int fstat ( FileDescriptor, Buffer)
```

```
int FileDescriptor;
```

```
struct stat *Buffer;
```

```
int statx (Path, Buffer, Length, Command)
```

```
char *Path;
```

```
struct stat *Buffer;
```

```
int Length;
```

```
int Command;
```

```
int fstatx (FileDescriptor, Buffer, Length, Command)
```

```
int FileDescriptor;
```

```
struct stat *Buffer;
```

```
int Length;
```

```
int Command;
```

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

```

int fullstat (Path, Command, Buffer)
struct fullstat *Buffer;
char *Path;
int Command;

int ffullstat (FileDescriptor, Command, Buffer)
struct fullstat *Buffer;
int FileDescriptor;
int Command;

int stat64 ( Path, Buffer)
const char *Path;
struct stat64 *Buffer;

int lstat64 (Path, Buffer)
const char *Path;
struct stat64 *Buffer;

int fstat64 ( FileDescriptor, Buffer)
int FileDescriptor;
struct stat64 *Buffer;

```

## Description

The **stat** subroutine obtains information about the file named by the *Path* parameter. Read, write, or execute permission for the named file is not required, but all directories listed in the path leading to the file must be searchable. The file information, which is a subset of the **stat** structure, is written to the area specified by the *Buffer* parameter.

The **lstat** subroutine obtains information about a file that is a symbolic link. The **lstat** subroutine returns information about the link, while the **stat** subroutine returns information about the file referenced by the link.

The **fstat** subroutine obtains information about the open file referenced by the *FileDescriptor* parameter. The **fstatx** subroutine obtains information about the open file referenced by the *FileDescriptor* parameter, as in the **fstat** subroutine.

The *st\_mode*, *st\_dev*, *st\_uid*, *st\_gid*, *st\_atime*, *st\_ctime*, and *st\_mtime* fields of the **stat** structure have meaningful values for all file types. The **statx**, **stat**, **lstat**, **fstatx**, **fstat**, **fullstat**, or **ffullstat** subroutine sets the *st\_nlink* field to a value equal to the number of links to the file.

The **statx** subroutine obtains a greater set of file information than the **stat** subroutine. The *Path* parameter is processed differently, depending on the contents of the *Command* parameter. The *Command* parameter provides the ability to collect information about symbolic links (as with the **lstat** subroutine) as well as information about mount points and hidden directories. The **statx** subroutine returns the amount of information specified by the *Length* parameter.

The **fullstat** and **ffullstat** subroutines are interfaces maintained for backward compatibility. With the exception of some field names, the **fullstat** structure is identical to the **stat** structure.

The **stat64**, **lstat64**, and **fstat64** subroutines are similar to the **stat**, **lstat**, **fstat** subroutines except that they return file information in a **stat64** structure instead of a **stat** structure. The information is identical except that the *st\_size* field is defined to be a 64-bit size. This allows **stat64**, **lstat64**, and **fstat64** to return file sizes which are greater than **OFF\_MAX** (2 gigabytes minus 1).

In the large file enabled programming environment, **stat** is redefined to be **stat64**, **lstat** is redefined to be **lstat64** and **fstat** is redefined to be **fstat64**.

## Parameters

*Path* Specifies the path name identifying the file. This name is interpreted differently depending on the interface used.

### *FileDescriptor*

Specifies the file descriptor identifying the open file.

*Buffer* Specifies a pointer to the **stat** structure in which information is returned. The **stat** structure is described in the **sys/stat.h** file.

### *Length*

Indicates the amount of information, in bytes, to be returned. Any value between 0 and the value returned by the **STATXSIZE** macro, inclusive, may be specified. The following macros may be used:

#### **STATSIZE**

Specifies the subset of the **stat** structure that is normally returned for a **stat** call.

#### **FULLSTATSIZE**

Specifies the subset of the **stat (fullstat)** structure that is normally returned for a **fullstat** call.

#### **STATXSIZE**

Specifies the complete **stat** structure. 0 specifies the complete **stat** structure, as if **STATXSIZE** had been specified.

### *Command*

Specifies a processing option. For the **statx** subroutine, the *Command* parameter determines how to interpret the path name provided, specifically, whether to retrieve information about a symbolic link, hidden directory, or mount point. Flags can be combined by logically ORing them together. The following options are possible values:

#### **STX\_LINK**

If the *Command* parameter specifies the **STX\_LINK** flag and the *Path* parameter is a path name that refers to a symbolic link, the **statx** subroutine returns information about the symbolic link. If the **STX\_LINK** flag is not specified, the **statx** subroutine returns information about the file to which the link refers.

If the *Command* parameter specifies the **STX\_LINK** flag and the *Path* value refers to a symbolic link, the *st\_mode* field of the returned **stat** structure indicates that the file is a symbolic link.

#### **STX\_HIDDEN**

If the *Command* parameter specifies the **STX\_HIDDEN** flag and the *Path* value is a path name that refers to a hidden directory, the **statx** subroutine returns information about the hidden directory. If the **STX\_HIDDEN** flag is not specified, the **statx** subroutine returns information about a subdirectory of the hidden directory.

If the *Command* parameter specifies the **STX\_HIDDEN** flag and *Path* refers to a hidden directory, the *st\_mode* field of the returned **stat** structure indicates that this is a hidden directory.

#### **STX\_MOUNT**

If the *Command* parameter specifies the **STX\_MOUNT** flag and the *Path* value is the name of a file or directory that has been mounted over, the **statx** subroutine returns information about the mounted-over file. If the **STX\_MOUNT** flag is not specified, the **statx** subroutine returns information about the mounted file or directory (the root directory of a virtual file system).

If the *Command* parameter specifies the **STX\_MOUNT** flag, the **FS\_MOUNT** bit in the *st\_flag* field of the returned **stat** structure is set if, and only if, this file is mounted over.

If the *Command* parameter does not specify the **STX\_MOUNT** flag, the **FS\_MOUNT** bit in the *st\_flag* field of the returned **stat** structure is set if, and only if, this file is the root directory of a virtual file system.

### STX\_NORMAL

If the *Command* parameter specifies the **STX\_NORMAL** flag, then no special processing is performed on the *Path* value. This option should be used when **STX\_LINK**, **STX\_HIDDEN**, and **STX\_MOUNT** flags are not desired.

For the **fstatx** subroutine, there are currently no special processing options. The only valid value for the *Command* parameter is the **STX\_NORMAL** flag.

For the **fullstat** and **ffullstat** subroutines, the *Command* parameter may specify the **FL\_STAT** flag, which is equivalent to the **STX\_NORMAL** flag, or the **FL\_NOFOLLOW** flag, which is equivalent to **STX\_LINK** flag.

### STX\_64

If the *Command* parameter specifies the **STX\_64** flag and the file size is greater than **OFF\_MAX**, then **statx** succeeds and returns the file size. Otherwise, **statx** fails and sets the **errno** to **Eoverflow**.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **stat**, **lstat**, **statx**, and **fullstat** subroutines are unsuccessful if one or more of the following are true:

<b>EACCES</b>	Search permission is denied for one component of the path prefix.
<b>ENAMETOOLONG</b>	The length of the path prefix exceeds the <b>PATH_MAX</b> flag value or a path name is longer than the <b>NAME_MAX</b> flag value while the <b>POSIX_NO_TRUNC</b> flag is in effect.
<b>ENOTDIR</b>	A component of the path prefix is not a directory.
<b>EFAULT</b>	Either the <i>Path</i> or the <i>Buffer</i> parameter points to a location outside of the allocated address space of the process.
<b>ENOENT</b>	The file named by the <i>Path</i> parameter does not exist.
<b>Eoverflow</b>	The size of the file is larger than can be represented in the <b>stat</b> structure pointed to by the <i>Buffer</i> parameter.

The **stat**, **lstat**, **statx**, and **fullstat** subroutines can be unsuccessful for other reasons. See "Base Operating System Error Codes for Services that Require Path-Name Resolution" for a list of additional errors.

The **fstat**, **fstatx**, and **ffullstat** subroutines fail if one or more of the following are true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter is not a valid file descriptor.
<b>EFAULT</b>	The <i>Buffer</i> parameter points to a location outside the allocated address space of the process.
<b>EIO</b>	An input/output (I/O) error occurred while reading from the file system.

The **statx** and **fstatx** subroutines are unsuccessful if one or more of the following are true:

<b>EINVAL</b>	The <i>Length</i> value is not between 0 and the value returned by the <b>STATSIZE</b> macro, inclusive.
<b>EINVAL</b>	The <i>Command</i> parameter contains an unacceptable value.

## Files

`/usr/include/sys/fullstat.h`

Contains the **fullstat** structure.

`/usr/include/sys/mode.h`

Defines values on behalf of the **stat.h** file.

## Related Information

The **chmod** subroutine, **chown** subroutine, **link** subroutine, **mknod** subroutine, **mount** (“vmount or mount Subroutine” on page 435) subroutine, **openx**, **open**, or **creat** subroutine, **pipe** subroutine, **symlink** (“symlink Subroutine” on page 308) subroutine, **vtimes** subroutine.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## strcat, strncat, strxfrm, strcpy, strncpy, or strdup Subroutine

### Purpose

Copies and appends strings in memory.

### Library

Standard C Library (**libc.a**)

### Syntax

**#include** <string.h>

```
char * strcat ( String1, String2)
char *String1;
const char *String2;
```

```
char * strncat (String1, String2, Number)
char *String1;
const char *String2;
size_t Number;
```

```
size_t strxfrm (String1, String2, Number)
char *String1;
const char *String2;
size_t Number;
```

```
char * strcpy (String1, String2)
char *String1;
const char *String2;
```

```
char * strncpy (String1, String2, Number)
char *String1;
const char *String2;
size_t Number;
```

```
char * strdup (String1)
const char *String1;
```

## Description

The **strcat**, **strncat**, **strxfrm**, **strcpy**, **strncpy**, and **strdup** subroutines copy and append strings in memory.

The *String1* and *String2* parameters point to strings. A string is an array of characters terminated by a null character. The **strcat**, **strncat**, **strcpy**, and **strncpy** subroutines all alter the string in the *String1* parameter. However, they do not check for overflow of the array to which the *String1* parameter points. String movement is performed on a character-by-character basis and starts at the left. Overlapping moves toward the left work as expected, but overlapping moves to the right may give unexpected results. All of these subroutines are declared in the **string.h** file.

The **strcat** subroutine adds a copy of the string pointed to by the *String2* parameter to the end of the string pointed to by the *String1* parameter. The **strcat** subroutine returns a pointer to the null-terminated result.

The **strncat** subroutine copies a number of bytes specified by the *Number* parameter from the *String2* parameter to the end of the string pointed to by the *String1* parameter. The subroutine stops copying before the end of the number of bytes specified by the *Number* parameter if it encounters a null character in the *String2* parameter's string. The **strncat** subroutine returns a pointer to the null-terminated result. The **strncat** subroutine returns the value of the *String1* parameter.

The **strxfrm** subroutine transforms the string pointed to by the *String2* parameter and places it in the array pointed to by the *String1* parameter. The **strxfrm** subroutine transforms the entire string if possible, but places no more than the number of bytes specified by the *Number* parameter in the array pointed to by the *String1* parameter. Consequently, if the *Number* parameter has a value of 0, the *String1* parameter can be a null pointer. The **strxfrm** subroutine returns the length of the transformed string, not including the terminating null byte. If the returned value is equal to or more than that of the *Number* parameter, the contents of the array pointed to by the *String1* parameter are indeterminable. If the number of bytes specified by the *Number* parameter is 0, the **strxfrm** subroutine returns the length required to store the transformed string, not including the terminating null byte. The **strxfrm** subroutine is determined by the **LC\_COLLATE** category.

The **strcpy** subroutine copies the string pointed to by the *String2* parameter to the character array pointed to by the *String1* parameter. Copying stops after the null character is copied. The **strcpy** subroutine returns the value of the *String1* parameter, if successful. Otherwise, a null pointer is returned.

The **strncpy** subroutine copies the number of bytes specified by the *Number* parameter from the string pointed to by the *String2* parameter to the character array pointed to by the *String1* parameter. If the *String2* parameter value is less than the specified number of characters, then the **strncpy** subroutine pads the *String1* parameter with trailing null characters to a number of bytes equaling the value of the *Number* parameter. If the *String2* parameter is exactly the specified number of characters or more, then only the number of characters specified by the *Number* parameter are copied and the result is not terminated with a null byte. The **strncpy** subroutine returns the value of the *String1* parameter.

The **strdup** subroutine returns a pointer to a new string, which is a duplicate of the string pointed to by the *String1* parameter. Space for the new string is obtained by using the **malloc** subroutine. A null pointer is returned if the new string cannot be created.

## Parameters

<i>Number</i>	Specifies the number of bytes in a string to be copied or transformed.
<i>String1</i>	Points to a string to which the specified data is copied or appended.
<i>String2</i>	Points to a string which contains the data to be copied, appended, or transformed.

## Error Codes

The **strcat**, **strncat**, **strxfrm**, **strcpy**, **strncpy**, and **strdup** subroutines fail if the following occurs:

**EFAULT** A string parameter is an invalid address.

In addition, the **strxfrm** subroutine fails if:

**EINVAL** A string parameter contains characters outside the domain of the collating sequence.

## Related Information

The **memccpy**, **memchr**, **memcmp**, **memcpy**, or **memmove** subroutine, **setlocale** (“setlocale Subroutine” on page 136) subroutine, **strcmp**, **strncmp**, **strcasecmp**, **strncasecmp**, or **strcoll** (“strcmp, strncmp, strcasecmp, strncasecmp, or strcoll Subroutine”) subroutine, **strlen**, **strchr**, **strchr**, **strpbrk**, **strspn**, **strcspn**, **strstr**, or **strtok** (“strlen, strchr, strchr, strpbrk, strspn, strcspn, strstr, strtok, or strsep Subroutine” on page 291) subroutine, **swab** (“swab Subroutine” on page 304) subroutine.

Subroutines, Example Programs, and Libraries and List of String Manipulation Services in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview, Multibyte and Wide Character String Collation Subroutines, and Multibyte and Wide Character String Comparison Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*

---

## strcmp, strncmp, strcasecmp, strncasecmp, or strcoll Subroutine

### Purpose

Compares strings in memory.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int strcmp ( String1, String2)
const char *String1, *String2;
```

```
int strncmp (String1, String2, Number)
const char *String1, *String2;
size_t Number;
```

```
int strcoll (String1, String2)
const char *String1, *String2;
```

```
#include <strings.h>
```

```
int strcasecmp (String1, String2)
const char *String1, *String2;
```

```
int strncasecmp (String1, String2, Number)
const char *String1, *String2;
size_t Number;
```

## Description

The **strcmp**, **strncmp**, **strcasecmp**, **strncasecmp**, and **strcoll** subroutines compare strings in memory.

The *String1* and *String2* parameters point to strings. A string is an array of characters terminated by a null character.

The **strcmp** subroutine performs a case-sensitive comparison of the string pointed to by the *String1* parameter and the string pointed to by the *String2* parameter, and analyzes the extended ASCII character set values of the characters in each string. The **strcmp** subroutine compares **unsigned char** data types. The **strcmp** subroutine then returns a value that is:

- Less than 0 if the value of string *String1* is lexicographically less than string *String2*.
- Equal to 0 if the value of string *String1* is lexicographically equal to string *String2*.
- Greater than 0 if the value of string *String1* is lexicographically greater than string *String2*.

The **strncmp** subroutine makes the same comparison as the **strcmp** subroutine, but compares up to the maximum number of pairs of bytes specified by the *Number* parameter.

The **strcasecmp** subroutine performs a character-by-character comparison similar to the **strcmp** subroutine. However, the **strcasecmp** subroutine is not case-sensitive. Uppercase and lowercase letters are mapped to the same character set value. The sum of the mapped character set values of each string is used to return a value that is:

- Less than 0 if the value of string *String1* is lexicographically less than string *String2*.
- Equal to 0 if the value of string *String1* is lexicographically equal to string *String2*.
- Greater than 0 if the value of string *String1* is lexicographically greater than string *String2*.

The **strncasecmp** subroutine makes the same comparison as the **strcasecmp** subroutine, but compares up to the maximum number of pairs of bytes specified by the *Number* parameter.

**Note:** Both the **strcasecmp** and **strncasecmp** subroutines only work with 7-bit ASCII characters.

The **strcoll** subroutine works the same as the **strcmp** subroutine, except that the comparison is based on a collating sequence determined by the **LC\_COLLATE** category. If the **strcmp** subroutine is used on transformed strings, it returns the same result as the **strcoll** subroutine for the corresponding untransformed strings.

## Parameters

<i>Number</i>	The number of bytes in a string to be examined.
<i>String1</i>	Points to a string which is compared.
<i>String2</i>	Points to a string which serves as the source for comparison.

## Error Codes

The **strcmp**, **strncmp**, **strcasecmp**, **strncasecmp**, and **strcoll** subroutines fail if the following occurs:

**EFAULT** A string parameter is an invalid address.

In addition, the **strcoll** subroutine fails if:

**EINVAL** A string parameter contains characters outside the domain of the collating sequence.

## Related Information

The **memccpy**, **memchr**, **memcmp**, **memcpy**, or **memmove** subroutine, **setlocale** (“setlocale Subroutine” on page 136) subroutine, **strcat**, **strncat**, **strxfrm**, **strcpy**, **strncpy**, or **strdup** (“strcat, strncat, strxfrm, strcpy, strncpy, or strdup Subroutine” on page 281) subroutine, **strlen**, **strchr**, **strchr**, **strpbrk**, **strspn**, **strcspn**, **strstr**, or **strtok** (“strlen, strchr, strchr, strpbrk, strspn, strcspn, strstr, strtok, or strsep Subroutine” on page 291) subroutine, **swab** (“swab Subroutine” on page 304) subroutine.

List of String Manipulation Subroutines and Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview, Multibyte and Wide Character String Collation Subroutines, and Multibyte and Wide Character String Comparison Subroutines *AIX 5L Version 5.2 National Language Support Guide and Reference*

---

## strerror Subroutine

### Purpose

Maps an error number to an error message string.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
char *strerror ( ErrorNumber )  
int ErrorNumber;
```

### Description

**Attention:** Do not use the **strerror** subroutine in a multithreaded environment.

The **strerror** subroutine maps the error number in the *ErrorNumber* parameter to the error message string. The **strerror** subroutine retrieves an error message based on the current value of the **LC\_MESSAGES** category. If the specified message catalog cannot be opened, the default message is returned. The returned message does not contain a new line (“\n”).

### Parameters

*ErrorNumber* Specifies the error number to be associated with the error message.

### Return Values

The **strerror** subroutine returns a pointer to the error message.

## Related Information

The **perror** subroutine.

The **clearerr** macro, **feof** macro, **ferror** macro, **fileno** macro.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## strfmon Subroutine

### Purpose

Formats monetary strings.

### Library

Standard C Library (**libc. a**)

### Syntax

```
#include <monetary.h>
```

```
ssize_t strfmon ( S, MaxSize, Format, ...)  
char *S;  
size_t MaxSize;  
const char *Format, ...;
```

### Description

The **strfmon** subroutine converts numeric values to monetary strings according to the specifications in the *Format* parameter. This parameter also contains numeric values to be converted. Characters are placed into the *S* array, as controlled by the *Format* parameter. The **LC\_MONETARY** category governs the format of the conversion.

The **strfmon** subroutine can be called multiple times by including additional **format** structures, as specified by the *Format* parameter.

The *Format* parameter specifies a character string that can contain plain characters and conversion specifications. Plain characters are copied to the output stream. Conversion specifications result in the fetching of zero or more arguments, which are converted and formatted.

If there are insufficient arguments for the *Format* parameter, the results are undefined. If arguments remain after the *Format* parameter is exhausted, the excess arguments are ignored.

A conversion specification consists of the following items in the following order: a % (percent sign), optional flags, optional field width, optional left precision, optional right precision, and a required conversion character that determines the conversion to be performed.

### Parameters

<i>S</i>	Contains the output of the <b>strfmon</b> subroutine.
<i>MaxSize</i>	Specifies the maximum number of bytes (including the null terminating byte) that may be placed in the <i>S</i> parameter.
<i>Format</i>	Contains characters and conversion specifications.

## Flags

One or more of the following flags can be specified to control the conversion:

- =f** An = (equal sign) followed by a single character that specifies the numeric fill character. 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 is specified.
- ^** Does not use grouping characters when formatting the currency amount. The default is to insert grouping characters if defined for the current locale.
- + or (** Determines the representation of positive and negative currency amounts. Only one of these flags may be specified. The locale's equivalent of + (plus sign) and - (negative sign) are used if + is specified. The locale's equivalent of enclosing negative amounts within parentheses is used if ( (left parenthesis) is specified. If neither flag is included, a default specified by the current locale is used.
- Left-justifies all fields (pads to the right). The default is right-justification.
- !** Suppresses the currency symbol from the output conversion.

## Field Width

- w** The decimal-digit string *w* specifies the minimum field width in which the result of the conversion is right-justified. If *-w* is specified, the result is left-justified. The default is a value of 0.

## Left Precision

- #n** A # (pound sign) followed by a decimal-digit string, *n*, specifies the maximum number of digits to be formatted to the left of the radix character. This option can be specified to keep formatted output from multiple calls to the **strfmon** subroutine aligned in the same columns. It can also be used to fill unused positions with a special character (for example, \$\*\*\*123.45). This option causes an amount to be formatted as if it has the number of digits specified by the *n* variable. If more than *n* digit positions are required, this option is ignored. Digit positions in excess of those required are filled with the numeric fill character set with the **=f** flag.

If defined for the current locale and not suppressed with the **^** flag, the subroutine inserts grouping characters before fill characters (if any). Grouping characters are not applied to fill characters, even if the fill character is a digit. In the example:

```
$0000001,234.56
```

grouping characters do not appear after the first or fourth 0 from the left.

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 equal in length.

## Right Precision

- .p** A . (period) followed by a decimal digit string, *p*, specifies the number of digits after the radix character. If the value of the *p* variable is 0, no radix character is used. If a right precision is not specified, a default specified by the current locale is use. The amount being formatted is rounded to the specified number of digits prior to formatting.

## Conversion Characters

- i** The double argument is formatted according to the current locale's international currency format; for example, in the U.S.: 1,234.56.
- n** The double argument is formatted according to the current locale's national currency format; for example, in the U.S.: \$1,234.56.
- %** No argument is converted; the conversion specification %% is replaced by a single %.

## Return Values

If successful, and if the number of resulting bytes (including the terminating null character) is not more than the number of bytes specified by the *MaxSize* parameter, the **strfmon** subroutine returns the number of bytes placed into the array pointed to by the *S* parameter (not including the terminating null byte). Otherwise, a value of -1 is returned and the contents of the *S* array are indeterminate.

## Error Codes

The **strfmon** subroutine may fail if the following is true:

**E2BIG** Conversion stopped due to lack of space in the buffer.

## Related Information

The **scanf** (“scanf, fscanf, sscanf, or wscanf Subroutine” on page 109) subroutine, **strftime** (“strftime Subroutine”) subroutine, **strptime** (“strptime Subroutine” on page 301) subroutine, **wcsftime** (“wcsftime Subroutine” on page 447) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and List of Time and Monetary Formatting Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## strftime Subroutine

### Purpose

Formats time and date.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
size_t strftime ( String, Length, Format, TmDate)
char *String;
size_t Length;
const char *Format;
const struct tm *TmDate;
```

### Description

The **strftime** subroutine converts the internal time and date specification of the **tm** structure, which is pointed to by the *TmDate* parameter, into a character string pointed to by the *String* parameter under the direction of the format string pointed to by the *Format* parameter. The actual values for the format specifiers are dependent on the current settings for the **LC\_TIME** category. The **tm** structure values may be assigned by the user or generated by the **localtime** or **gmtime** subroutine. The resulting string is similar to the result of the **printf** *Format* parameter, and is placed in the memory location addressed by the *String* parameter. The maximum length of the string is determined by the *Length* parameter and terminates with a null character.

Many conversion specifications are the same as those used by the **date** command. The interpretation of some conversion specifications is dependent on the current locale of the process.

The *Format* parameter is a character string containing two types of objects: plain characters that are simply placed in the output string, and conversion specifications that convert information from the *TmDate* parameter into readable form in the output string. Each conversion specification is a sequence of this form:

% type

- A % (percent sign) introduces a conversion specification.
- The type of conversion is specified by one or two conversion characters. The characters and their meanings are:

<b>%a</b>	Represents the locale's abbreviated weekday name (for example, Sun) defined by the <b>abday</b> statement in the <b>LC_TIME</b> category.
<b>%A</b>	Represents the locale's full weekday name (for example, Sunday) defined by the <b>day</b> statement in the <b>LC_TIME</b> category.
<b>%b</b>	Represents the locale's abbreviated month name (for example, Jan) defined by the <b>abmon</b> statement in the <b>LC_TIME</b> category.
<b>%B</b>	Represents the locale's full month name (for example, January) defined by the <b>mon</b> statement in the <b>LC_TIME</b> category.
<b>%c</b>	Represents the locale's date and time format defined by the <b>d_t_fmt</b> statement in the <b>LC_TIME</b> category.
<b>%C</b>	Represents the century number (the year divided by 100 and truncated to an integer) as a decimal number (00 through 99).
<b>%d</b>	Represents the day of the month as a decimal number (01 to 31).
<b>%D</b>	Represents the date in <b>%m/%d/%y</b> format (for example, 01/31/91).
<b>%e</b>	Represents the day of the month as a decimal number (01 to 31). The <b>%e</b> field descriptor uses a two-digit field. If the day of the month is not a two-digit number, the leading digit is filled with a space character.
<b>%E</b>	Represents the locale's combined alternate era year and name, respectively, in <b>%o %N</b> format.
<b>%h</b>	Represents the locale's abbreviated month name (for example, Jan) defined by the <b>abmon</b> statement in the <b>LC_TIME</b> category. This field descriptor is a synonym for the <b>%b</b> field descriptor.
<b>%H</b>	Represents the 24-hour-clock hour as a decimal number (00 to 23).
<b>%I</b>	Represents the 12-hour-clock hour as a decimal number (01 to 12).
<b>%j</b>	Represents the day of the year as a decimal number (001 to 366).
<b>%m</b>	Represents the month of the year as a decimal number (01 to 12).
<b>%M</b>	Represents the minutes of the hour as a decimal number (00 to 59).
<b>%n</b>	Specifies a new-line character.
<b>%N</b>	Represents the locale's alternate era name.
<b>%o</b>	Represents the alternate era year.
<b>%p</b>	Represents the locale's a.m. or p.m. string defined by the <b>am_pm</b> statement in the <b>LC_TIME</b> category.
<b>%r</b>	Represents 12-hour clock time with a.m./p.m. notation as defined by the <b>t_fmt_ampm</b> statement. The usual format is <b>%I:%M:%S %p</b> .
<b>%R</b>	Represents 24-hour clock time in <b>%H:%M</b> format.
<b>%S</b>	Represents the seconds of the minute as a decimal number (00 to 59).
<b>%t</b>	Specifies a tab character.
<b>%T</b>	Represents 24-hour-clock time in the format <b>%H:%M:%S</b> (for example, 16:55:15).
<b>%u</b>	Represents the weekday as a decimal number (1 to 7). Monday or its equivalent is considered the first day of the week for calculating the value of this field descriptor.
<b>%U</b>	Represents the week of the year as a decimal number (00 to 53). Sunday, or its equivalent as defined by the <b>day</b> statement in the <b>LC_TIME</b> category, is considered the first day of the week for calculating the value of this field descriptor.
<b>%V</b>	Represents the week number of the year (with Monday as the first day of the week) as a decimal number (01 to 53). If the week containing January 1 has four or more days in the new year, then it is considered week 1; otherwise, it is considered week 53 of the previous year, and the next week is week 1 of the new year.
<b>%w</b>	Represents the day of the week as a decimal number (0 to 6). Sunday, or its equivalent as defined by the <b>day</b> statement, is considered as 0 for calculating the value of this field descriptor.
<b>%W</b>	Represents the week of the year as a decimal number (00 to 53). Monday, or its equivalent as defined by the <b>day</b> statement, is considered the first day of the week for calculating the value of this field descriptor.
<b>%x</b>	Represents the locale's date format as defined by the <b>d_fmt</b> statement.
<b>%X</b>	Represents the locale's time format as defined by the <b>t_fmt</b> statement.

- %y** Represents the year of the century.  
**Note:** When the environment variable **XPG\_TIME\_FMT=ON**, **%y** is the year within the 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 2000 to 2068, inclusive.
- %Y** Represents the year as a decimal number (for example, 1989).
- %Z** Represents the time-zone name if one can be determined (for example, EST). No characters are displayed if a time zone cannot be determined.
- %%** Specifies a % (percent sign).

Some conversion specifiers can be modified by the **E** or **O** modifier characters to indicate that an alternative format or specification should be used. If the alternative format or specification does not exist for the current locale, the behavior will be the same as with the unmodified conversion specification. The following modified conversion specifiers are supported:

- %Ec** Represents the locale's alternative appropriate date and time as defined by the **era\_d\_t\_fmt** statement.
- %EC** Represents the name of the base year (or other time period) in the locale's alternative form as defined by the **era** statement under the **era\_name** category of the current era.
- %Ex** Represents the locale's alternative date as defined by the **era\_d\_fmt** statement.
- %EX** Represents the locale's alternative time as defined by the **era\_t\_fmt** statement.
- %Ey** Represents the offset from the **%EC** modified conversion specifier (year only) in the locale's alternative form.
- %EY** Represents the full alternative-year form.
- %Od** Represents the day of the month, using the locale's alternative numeric symbols, filled as needed with leading 0's if an alternative symbol for 0 exists. If an alternative symbol for 0 does not exist, the **%Od** modified conversion specifier uses leading space characters.
- %Oe** Represents the day of the month, using the locale's alternative numeric symbols, filled as needed with leading 0's if an alternative symbol for 0 exists. If an alternative symbol for 0 does not exist, the **%Oe** modified conversion specifier uses leading space characters.
- %OH** Represents the hour in 24-hour clock time, using the locale's alternative numeric symbols.
- %OI** Represents the hour in 12-hour clock time, using the locale's alternative numeric symbols.
- %Om** Represents the month, using the locale's alternative numeric symbols.
- %OM** Represents the minutes, using the locale's alternative numeric symbols.
- %OS** Represents the seconds, using the locale's alternative numeric symbols.
- %Ou** Represents the weekday as a number using the locale's alternative numeric symbols.
- %OU** Represents the week number of the year, using the locale's alternative numeric symbols. Sunday is considered the first day of the week. Use the rules corresponding to the **%U** conversion specifier.
- %OV** Represents the week number of the year (Monday as the first day of the week, rules corresponding to **%V**) using the locale's alternative numeric symbols.
- %Ow** Represents the number of the weekday (with Sunday equal to 0), using the locale's alternative numeric symbols.
- %OW** Represents the week number of the year using the locale's alternative numeric symbols. Monday is considered the first day of the week. Use the rules corresponding to the **%W** conversion specifier.
- %Oy** Represents the year (offset from **%C**) using the locale's alternative numeric symbols.

## Parameters

- String* Points to the string to hold the formatted time.
- Length* Specifies the maximum length of the string pointed to by the *String* parameter.
- Format* Points to the format character string.
- TmDate* Points to the time structure that is to be converted.

## Return Values

If the total number of resulting bytes, including the terminating null byte, is not more than the *Length* value, the **strftime** subroutine returns the number of bytes placed into the array pointed to by the *String* parameter, not including the terminating null byte. Otherwise, a value of 0 is returned and the contents of the array are indeterminate.

## Related Information

The **localtime** subroutine, **gmtime** subroutine, **mbstowcs** subroutine, **printf** subroutine, **strfmon** (“strfmon Subroutine” on page 286) subroutine, **strptime** (“strptime Subroutine” on page 301) subroutine, **wcsftime** (“wcsftime Subroutine” on page 447) subroutine.

The **date** command.

LC\_TIME Category for the Locale Definition Source File Format in *AIX 5L Version 5.2 Files Reference*.

List of Time Data Manipulation Services in *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

, Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## strlen, strchr, strchr, strpbrk, strspn, strcspn, strstr, strtok, or strsep Subroutine

### Purpose

Determines the size, location, and existence of strings in memory.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
size_t strlen (String)
const char *String;
```

```
char *strchr (String, Character)
const char *String;
int Character;
```

```
char *strchr (String, Character)
const char *String;
int Character;
```

```
char *strpbrk (String1, String2)
const char *String1, *String2;
```

```
size_t strspn (String1, String2)
const char *String1, *String2;
```

```
size_t strcspn (String1, String2)
const char *String1, *String2;
```

```
char *strstr (String1, String2)
const char *String1, *String2;
```

```
char *strtok (String1, String2)
char *String1;
const char *String2;
```

```
char *strsep (String1, String2)
char **String1;
const char *String2;
```

```
char *index (String, Character)
const char *String;
int Character;
```

```
char *rindex (String, Character)
const char *String;
int Character;
```

## Description

**Attention:** Do not use the **strtok** subroutine in a multithreaded environment. Use the **strtok\_r** subroutine instead.

The **strlen**, **strchr**, **strrchr**, **strpbrk**, **strspn**, **strcspn**, **strstr**, and **strtok** subroutines determine such values as size, location, and the existence of strings in memory.

The *String1*, *String2*, and *String* parameters point to strings. A string is an array of characters terminated by a null character.

The **strlen** subroutine returns the number of bytes in the string pointed to by the *String* parameter, not including the terminating null bytes.

The **strchr** subroutine returns a pointer to the first occurrence of the character specified by the *Character* (converted to an unsigned character) parameter in the string pointed to by the *String* parameter. A null pointer is returned if the character does not occur in the string. The null byte that terminates a string is considered to be part of the string.

The **strrchr** subroutine returns a pointer to the last occurrence of the character specified by the *Character* (converted to a character) parameter in the string pointed to by the *String* parameter. A null pointer is returned if the character does not occur in the string. The null byte that terminates a string is considered to be part of the string.

The **strpbrk** subroutine returns a pointer to the first occurrence in the string pointed to by the *String1* parameter of any bytes from the string pointed to by the *String2* parameter. A null pointer is returned if no bytes match.

The **strspn** subroutine returns the length of the initial segment of the string pointed to by the *String1* parameter, which consists entirely of bytes from the string pointed to by the *String2* parameter.

The **strcspn** subroutine returns the length of the initial segment of the string pointed to by the *String1* parameter, which consists entirely of bytes *not* from the string pointed to by the *String2* parameter.

The **strstr** subroutine finds the first occurrence in the string pointed to by the *String1* parameter of the sequence of bytes specified by the string pointed to by the *String2* parameter (excluding the terminating null character). It returns a pointer to the string found in the *String1* parameter, or a null pointer if the string was not found. If the *String2* parameter points to a string of 0 length, the **strstr** subroutine returns the value of the *String1* parameter.

The **strtok** subroutine breaks the string pointed to by the *String1* parameter into a sequence of tokens, each of which is delimited by a byte from the string pointed to by the *String2* parameter. The first call in the sequence takes the *String1* parameter as its first argument and is followed by calls that take a null pointer as their first argument. The separator string pointed to by the *String2* parameter may be different from call to call.

The first call in the sequence searches the *String1* parameter for the first byte that is not contained in the current separator string pointed to by the *String2* parameter. If no such byte is found, no tokens exist in the string pointed to by the *String1* parameter, and a null pointer is returned. If such a byte is found, it is the start of the first token.

The **strtok** subroutine then searches from the first token 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 the *String1* parameter, and subsequent searches for a token return a null pointer. If such a byte is found, the **strtok** subroutine overwrites it with a null byte, which terminates the current token. The **strtok** subroutine saves a pointer to the following byte, from which the next search for a token will start. The subroutine returns a pointer to the first byte of the token.

Each subsequent call with a null pointer as the value of the first argument starts searching from the saved pointer, using it as the first token. Otherwise, the subroutine's behavior does not change.

The **strsep** subroutine returns the next token from the string *String1* which is delimited by *String2*. The token is terminated with a `\0` character and *String1* is updated to point past the token. The **strsep** subroutine returns a pointer to the token, or NULL if *String2* is not found in *String1*.

The **index**, **rindex** and **strsep** subroutines are included for compatibility with BSD and are not part of the ANSI C Library. The **index** subroutine is implemented as a call to the **strchr** subroutine. The **rindex** subroutine is implemented as a call to the **strrchr** subroutine.

## Parameters

<i>Character</i>	Specifies a character for which to return a pointer.
<i>String</i>	Points to a string from which data is returned.
<i>String1</i>	Points to a string from which an operation returns results.
<i>String2</i>	Points to a string which contains source for an operation.

## Error Codes

The **strlen**, **strchr**, **strrchr**, **strpbrk**, **strspn**, **strcspn**, **strstr**, and **strtok** subroutines fail if the following occurs:

**EFAULT** A string parameter is an invalid address.

## Related Information

The “setlocale Subroutine” on page 136, “strcat, strncat, strxfrm, strcpy, strncpy, or strdup Subroutine” on page 281, “strcmp, strncmp, strcasecmp, strncasecmp, or strcoll Subroutine” on page 283, “strtok\_r Subroutine” on page 298, and “swab Subroutine” on page 304.

The **memccpy**, **memchr**, **memcmp**, **memcpy**, or **memmove** subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*

List of String Manipulation Services and Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

## strncollen Subroutine

### Purpose

Returns the number of collation values for a given string.

### Library

Standard C Library (**libc.a**)

### Syntax

```
include <string.h>
```

```
int strncollen ( String, Number )  
const char *String;  
const int Number;
```

### Description

The **strncollen** subroutine returns the number of collation values for a given string pointed to by the *String* parameter. The count of collation values is terminated when either a null character is encountered or when the number of bytes indicated by the *Number* parameter have been examined.

The collation values are set by the **setlocale** subroutine for the **LC\_COLLATE** category. For example, if the locale is set to **Es\_ES** (Spanish spoken in Spain) for the **LC\_COLLATE** category, where `ch` has one collation value, then **strncollen** ('abchd', 5) returns 4.

In German, the <Sharp-S> character has two collation values, so substituting the <Sharp-S> character for B in the following example, **strncollen** ('straBa', 6) returns 7.

If a character has no collation value, its collation length is 0.

### Parameters

*Number*     The number of bytes in a string to be examined.  
*String*     Pointer to a string to be examined for collation value.

### Return Values

Upon successful completion, the **strncollen** subroutine returns the collation value for a given string, pointed to by the *String* parameter.

### Related Information

The **setlocale** ("setlocale Subroutine" on page 136) subroutine, **strcat**, **strncat**, **strxfrm**, **strcpy**, **strncpy**, or **strdup** ("strcat, strncat, strxfrm, strcpy, strncpy, or strdup Subroutine" on page 281) subroutine, **strcmp**, **strncmp**, **strcasestr**, **strncasestr**, or **strcoll** ("strcmp, strncmp, strcasestr, strncasestr, or strcoll Subroutine" on page 283) subroutine, **strlen**, **strchr**, **strrchr**, **strpbrk**, **strspn**, **strcspn**, **strstr**, or **strtok** ("strlen, strchr, strrchr, strpbrk, strspn, strcspn, strstr, strtok, or strsep Subroutine" on page 291) subroutine.

---

## strtouf, strtod, or strtold Subroutine

### Purpose

Converts a string to a double-precision number.

### Syntax

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

```
float strtouf (nptr, endptr)  
const char *restrict nptr;  
char **restrict endptr;
```

```
double strtod ( nptr, endptr)  
const char *nptr  
char**endptr;
```

```
long double strtold (nptr, endptr)  
const char *restrict nptr;  
char **restrict endptr;
```

### Description

The **strtouf**, **strtod**, and **strtold** subroutines convert the initial portion of the string pointed to by *nptr* to **double**, **float**, and **long double** representation, respectively. First, they decompose the input string into three parts:

- An initial, possibly empty, sequence of white-space characters (as specified by `isspace()`).
- A subject sequence interpreted as a floating-point constant or representing infinity or NaN.
- A final string of one or more unrecognized characters, including the terminating null byte of the input string.

Then, they attempt to convert the subject sequence to a floating-point number, and return the result.

The expected form of the subject sequence is an optional plus or minus sign, and one of the following:

- A non-empty sequence of decimal digits optionally containing a radix character, and an optional exponent part
- A 0x or 0X, and a non-empty sequence of hexadecimal digits optionally containing a radix character, and an optional binary exponent part
- One of INF or INFINITY, ignoring case
- One of NAN or NAN(*n-char-sequence* *opt*), ignoring case in the NAN part, where:

```
n-char-sequence:  
  digit  
  nondigit  
  n-char-sequence digit  
  n-char-sequence nondigit
```

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) are interpreted as a floating constant of the C language, except that the radix character is used in place of a period, and 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 shall be interpreted as an infinity, if representable in the return type, or else as if it were a floating constant that is too large for the range of the return type. A character sequence NAN or NAN(*n-char-sequence*<sub>opt</sub>) is interpreted as a quiet NaN, if supported in the return type, or else as if it were a subject sequence part that does not have the expected form. The meaning of the *n*-char sequences is implementation-defined. A pointer to the final string is stored in the object pointed to by the *endptr* parameter, provided that the *endptr* parameter is not a null pointer.

If the subject sequence has the hexadecimal form, the value resulting from the conversion is correctly rounded.

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.

In other than the C or POSIX locales, other implementation-defined subject sequences may be accepted.

If the subject sequence is empty or does not have the expected form, no conversion shall be performed; the value of **str** is stored in the object pointed to by *endptr*, provided that *endptr* is not a null pointer.

The **strtod** subroutine does not change the setting of the **errno** global variable if successful.

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, call the **strtof** or **strtold** subroutine, then check **errno**.

## Parameters

<i>nptr</i>	Specifies the string to be converted.
<i>endptr</i>	Points to the final string.

## Return Values

Upon successful completion, the **strtof** and **strtold** subroutines return the converted value. If no conversion could be performed, 0 is returned, and the **errno** global variable may be set to EINVAL.

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), and **errno** is set to ERANGE.

If the correct value would cause an underflow, a value whose magnitude is no greater than the smallest normalized positive number in the return type is returned and the **errno** global variable is set to ERANGE.

## Error Codes

**Note:** Because a value of 0 can indicate either an error or a valid result, an application that checks for errors with the **strtod**, **strtof**, and **strtold** subroutines should set the **errno** global variable equal to 0 prior to the subroutine call. The application can check the **errno** global variable after the subroutine call.

If the string pointed to by *NumberPointer* is empty or begins with an unrecognized character, a value of 0 is returned for the **strtod**, **strtof**, and **strtold** subroutines.

If the conversion cannot be performed, a value of 0 is returned, and the **errno** global variable is set to indicate the error.

If the conversion causes an overflow (that is, the value is outside the range of representable values), **+/- HUGE\_VAL** is returned with the sign indicating the direction of the overflow, and the **errno** global variable is set to **ERANGE**.

If the conversion would cause an underflow, a properly signed value of 0 is returned and the **errno** global variable is set to **ERANGE**.

For the **strtod**, **strtof**, and **strtold** subroutines, if the value of the *EndPointer* parameter is not (**char\*\***) NULL, a pointer to the character that stopped the subroutine is stored in *\*EndPointer*. If a floating-point value cannot be formed, *\*EndPointer* is set to *NumberPointer*.

The **strtof** subroutine has only one rounding error. (If the **strtod** subroutine is used to create a double-precision floating-point number and then that double-precision number is converted to a floating-point number, two rounding errors could occur.)

## Related Information

“scanf, fscanf, sscanf, or wscanf Subroutine” on page 109, “setlocale Subroutine” on page 136, and “strtol, strtoul, strtoll, strtoull, or atoi Subroutine” on page 299.

cctype, isalpha, isupper, islower, isdigit, isxdigit, isalnum, isspace, ispunct, isprint, isgraph, iscntrl, or isascii Subroutines and localeconv Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

---

## strtoimax or strtoumax Subroutine

### Purpose

Converts string to integer type.

### Syntax

```
#include <inttypes.h>
```

```
intmax_t strtoimax (nptr, endptr, base)
const char *restrict nptr;
char **restrict endptr;
int base;
```

```
uintmax_t strtoumax (nptr, endptr, base)
const char *restrict nptr;
char **restrict endptr;
int base;
```

### Description

The **strtoimax** and **strtoumax** subroutines are equivalent to the **strtol**, **strtoll**, **strtoul**, and **strtoull** subroutines, except that the initial portion of the string shall be converted to **intmax\_t** and **uintmax\_t** representation, respectively.

### Parameters

<i>nptr</i>	Points to the string to be converted.
<i>endptr</i>	Points to the object where the final string is stored.
<i>base</i>	Determines the value of the integer represented in some radix.

## Return Values

The **strtoimax** and **strtoumax** subroutines return the converted value, if any.

If no conversion could be performed, zero is returned.

If the correct value is outside the range of representable values, **{INTMAX\_MAX}**, **{INTMAX\_MIN}**, or **{UINTMAX\_MAX}** is returned (according to the return type and sign of the value, if any), and the **errno** global variable is set to **ERANGE**.

## Related Information

The “**strtol**, **strtoul**, **strtoll**, **strtoull**, or **atoi** Subroutine” on page 299.

inttypes.h in *AIX 5L Version 5.2 Files Reference*.

---

## strtok\_r Subroutine

### Purpose

Breaks a string into a sequence of tokens.

### Libraries

Thread-Safe C Library (**libc\_r.a**)

### Syntax

```
#include<string.h>
char *strtok_r (String, Separators, Pointer);
char *String;
const char *Separators;
char **Pointer;
```

### Description

**Note:** The **strtok\_r** subroutine is used in a multithreaded environment.

The **strtok\_r** subroutine breaks the string pointed to by the *String* parameter into a sequence of tokens, each of which is delimited by a byte from the string pointed to by the *Separators* parameter. The *Pointer* parameter holds the information necessary for the **strtok\_r** subroutine to perform scanning on the *String* parameter. In the first call to the **strtok\_r** subroutine, the value passed as the *Pointer* parameter is ignored.

The first call in the sequence searches the *String* parameter for the first byte that is not contained in the current separator string pointed to by the *Separators* parameter. If no such byte is found, no tokens exist in the *String* parameter, and a null pointer is returned. If such a byte is found, it is the start of the first token. The **strtok\_r** subroutine also updates the *Pointer* parameter with the starting address of the token following the first occurrence of the *Separators* parameter.

In subsequent calls, a null pointer should be passed as the first parameter to the **strtok\_r** subroutine instead of the *String* parameter. Each subsequent call with a null pointer as the value of the first argument starts searching from the *Pointer* parameter, using it as the first token. Otherwise, the subroutine's behavior does not change. The **strtok\_r** subroutine would return successive tokens until no tokens remain. The *Separators* parameter may be different from one call to another.

### Parameters

*String* Points to a string from which an operation returns results.

Separators  
Pointer

Points to a string which contains source for an operation.  
Points to a user provided pointer.

## Error Codes

The `strtok_r` subroutine fails if the following occurs:

**EFAULT**

A *String* parameter is an invalid address.

## Related Information

The “`strlen`, `strchr`, `strchr`, `strpbrk`, `strspn`, `strcspn`, `strstr`, `strtok`, or `strsep` Subroutine” on page 291.

Writing Reentrant and Thread-Safe Code in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## strtol, strtoul, strtoll, strtoull, or atoi Subroutine

### Purpose

Converts a string to a signed or unsigned long integer or long long integer.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
long strtol (String, EndPointer, Base)
```

```
const char *String;
```

```
char **EndPointer;
```

```
int Base;
```

```
unsigned long strtoul (String, EndPointer, Base)
```

```
const char *String;
```

```
char **EndPointer;
```

```
int Base;
```

```
long long int strtoll (String, EndPointer, Base)
```

```
char *String, **EndPointer;
```

```
int Base;
```

```
long long int strtoull (String, EndPointer, Base)
```

```
char *String, **EndPointer;
```

```
int Base;
```

```
int atoi (String)
```

```
const char *String;
```

### Description

The `strtol` subroutine returns a long integer whose value is represented by the character string to which the *String* parameter points. The `strtol` subroutine scans the string up to the first character that is inconsistent with the *Base* parameter. Leading white-space characters are ignored, and an optional sign may precede the digits.

The `strtoul` subroutine provides the same functions but returns an unsigned long integer.

The **strtol** and **strtoull** subroutines provide the same functions but return long long integers.

The **atoi** subroutine is equivalent to the **strtol** subroutine where the value of the *EndPoint* parameter is a null pointer and the *Base* parameter is a value of 10.

If the value of the *EndPoint* parameter is not null, then a pointer to the character that ended the scan is stored in *EndPoint*. If an integer cannot be formed, the value of the *EndPoint* parameter is set to that of the *String* parameter.

If the *Base* parameter is a value between 2 and 36, the subject sequence's expected form is a sequence of letters and digits representing an integer whose radix is specified by the *Base* parameter. This sequence is optionally preceded by a + (positive) or - (negative) sign. 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 the *Base* parameter are permitted. If the *Base* parameter has a value of 16, the characters 0x or 0X optionally precede the sequence of letters and digits, following the + (positive) or - (negative) sign if present.

If the value of the *Base* parameter is 0, the string determines the base. Thus, after an optional leading sign, a leading 0 indicates octal conversion, and a leading 0x or 0X indicates hexadecimal conversion. The default is to use decimal conversion.

## Parameters

<i>String</i>	Points to the character string to be converted.
<i>EndPoint</i>	Points to a character string that contains the first character not converted.
<i>Base</i>	Specifies the base to use for the conversion.

## Return Values

Upon successful completion, the **strtol**, **strtoul**, **strtol**, and **strtoull** subroutines return the converted value. If no conversion could be performed, 0 is returned, and the **errno** global variable is set to indicate the error. If the correct value is outside the range of representable values, the **strtol** subroutine returns a value of **LONG\_MAX** or **LONG\_MIN** according to the sign of the value, while the **strtoul** subroutine returns a value of **ULONG\_MAX**. The **strtol** subroutine returns a value of **LLONG\_MAX** or **LLONG\_MIN**, according to the sign of the value. The **strtoul** subroutine returns a value of **ULONG\_MAX**, and the **strtoull** subroutine returns a value of **ULLONG\_MAX**.

## Error Codes

The **strtol** and **strtoul** subroutines return the following error codes:

<b>ERANGE</b>	The correct value of the converted number causes underflow or overflow.
<b>EINVAL</b>	The value of the <i>Base</i> parameter is not valid.

## Related Information

The **atof**, **atoff**, **strtod**, or **strtof** subroutine, **scanf**, **fscanf**, **sscanf**, or **wsscanf** ("scanf, fscanf, sscanf, or wsscanf Subroutine" on page 109) subroutine, **setlocale** ("setlocale Subroutine" on page 136) subroutine, **wstrtod** or **watof** ("wstrtod or watof Subroutine" on page 514) subroutine, **wstrtol**, **watol**, or **watoi** ("wstrtol, watol, or watoi Subroutine" on page 515) subroutine.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## strptime Subroutine

### Purpose

Converts a character string to a time value.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
char *strptime ( Buf, Format, Tm)  
const char *Buf, *Format;  
struct tm *Tm;
```

### Description

The **strptime** subroutine converts the characters in the *Buf* parameter to values that are stored in the *Tm* structure, using the format specified by the *Format* parameter.

### Parameters

*Buf*            Contains the character string to be converted by the **strptime** subroutine.

*Format*        Contains format specifiers for the **strptime** subroutine. The *Format* parameter contains 0 or more specifiers. Each specifier is composed of one of the following elements:

- One or more white-space characters
- An ordinary character (neither % [percent sign] nor a white-space character)
- A format specifier

**Note:** If more than one format specifier is present, they must be separated by white space or a non-percent/non-alphanumeric character. If the separator between format specifiers is other than white space, the *Buf* string should hold the same separator at the corresponding locations.

The **LC\_TIME** category defines the locale values for the format specifiers. The following format specifiers are supported:

**%a**    Represents the weekday name, either abbreviated as specified by the **abday** statement or full as specified by the **day** statement.

**%A**    Represents the weekday name, either abbreviated as specified by the **abday** statement or full as specified by the **day** statement.

**%b**    Represents the month name, either abbreviated as specified by the **abmon** statement or full as specified by the **month** statement.

**%B**    Represents the month name, either abbreviated as specified by the **abmon** statement or full as specified by the **month** statement.

**%c**    Represents the date and time format defined by the **d\_t\_fmt** statement in the **LC\_TIME** category.

**%C**    Represents the century number (0 through 99); leading zeros are permitted but not required.

**%d**    Represents the day of the month as a decimal number (01 to 31).

**%D**    Represents the date in **%m/%d/%y** format (for example, 01/31/91).

**%e**    Represents the day of the month as a decimal number (01 to 31).

**%E**    Represents the combined alternate era year and name, respectively, in **%o %N** format.

**%h**    Represents the month name, either abbreviated as specified by the **abmon** statement or full as specified by the **month** statement.

**%H**    Represents the 24-hour-clock hour as a decimal number (00 to 23).

- %l** Represents the 12-hour-clock hour as a decimal number (01 to 12).
- %j** Represents the day of the year as a decimal number (001 to 366).
- %m** Represents the month of the year as a decimal number (01 to 12).
- %M** Represents the minutes of the hour as a decimal number (00 to 59).
- %n** Represents any white space.
- %N** Represents the alternate era name.
- %o** Represents the alternate era year.
- %p** Represents the a.m. or p.m. string defined by the **am\_pm** statement in the **LC\_TIME** category.
- %r** Represents 12-hour-clock time with a.m./p.m. notation as defined by the **t\_fmt\_ampm** statement, usually in the format **!:%M:%S %p**.
- %S** Represents the seconds of the minute as a decimal number (00 to 61). The decimal number range of 00 to 61 provides for leap seconds.
- %t** Represents any white space.
- %T** Represents 24-hour-clock time in the format **%H:%M:%S** (for example, 16:55:15).
- %U** Represents the week of the year as a decimal number (00 to 53). Sunday, or its equivalent as defined by the **day** statement, is considered the first day of the week for calculating the value of this field descriptor.
- %w** Represents the day of the week as a decimal number (0 to 6). Sunday, or its equivalent as defined by the **day** statement in the **LC\_TIME** category, is considered to be 0 for calculating the value of this field descriptor.
- %W** Represents the week of the year as a decimal number (00 to 53). Monday, or its equivalent as defined by the **day** statement in the **LC\_TIME** category, is considered the first day of the week for calculating the value of this field descriptor.
- %x** Represents the date format defined by the **d\_fmt** statement in the **LC\_TIME** category.
- %X** Represents the time format defined by the **t\_fmt** statement in the **LC\_TIME** category.
- %y** Represents the year within century.  
**Note:** When the environment variable **XPG\_TIME\_FMT=ON**, **%y** is the year within the 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 2000 to 2068, inclusive.
- %Y** Represents the year as a decimal number (for example, 1989).
- %Z** Represents the time-zone name, if one can be determined (for example, EST). No characters are displayed if a time zone cannot be determined.
- %%** Specifies a % (percent sign) character.

Some format specifiers can be modified by the **E** and **O** modifier characters to indicate an alternative format or specification. If the alternative format or specification does not exist in the current locale, the behavior will be as if the unmodified format specifier were used. The following modified format specifiers are supported:

- %Ec** Represents the locale's alternative appropriate date and time as defined by the **era\_d\_t\_fmt** statement.
- %EC** Represents the base year (or other time period) in the locale's alternative form as defined by the **era** statement under the **era\_name** category of the current era.
- %Ex** Represents the alternative date as defined by the **era\_d\_fmt** statement.
- %EX** Represents the locale's alternative time as defined by the **era\_t\_fmt** statement.
- %Ey** Represents the offset from the **%EC** format specifier (year only) in the locale's alternative form.
- %EY** Represents the full alternative-year format.
- %Od** Represents the month using the locale's alternative numeric symbols. Leading 0's are permitted but not required.
- %Oe** Represents the month using the locale's alternative numeric symbols. Leading 0's are permitted but not required.
- %OH** Represents the hour in 24-hour-clock time using the locale's alternative numeric symbols.
- %OI** Represents the hour in 12-hour-clock time using the locale's alternative numeric symbols.
- %Om** Represents the month using the locale's alternative numeric symbols.
- %OM** Represents the minutes using the locale's alternative numeric symbols.
- %OS** Represents the seconds using the locale's alternative numeric symbols.
- %OU** Represents the week number of the year using the locale's alternative numeric symbols. Sunday is considered the first day of the week. Use the rules corresponding to the **%U** format specifier.

- %Ow** Represents the day of the week using the locale's alternative numeric symbols. Sunday is considered the first day of the week.
- %OW** Represents the week number of the year using the locale's alternative numeric symbols. Monday is considered the first day of the week. Use the rules corresponding to the **%W** format specifier.
- %Oy** Represents the year (offset from **%C**) using the locale's alternative numeric symbols.

A format specification consisting of white-space characters is performed by reading input until the first nonwhite-space character (which is not read) or up to no more characters can be read.

A format specification consisting of an ordinary character is performed by reading the next character from the *Buf* parameter. If this character differs from the character comprising the directive, the directive fails and the differing character and any characters following it remain unread. Case is ignored when matching *Buf* items, such as month or weekday names.

A series of directives composed of **%n** format specifiers, **%t** format specifiers, white-space characters, or any combination of the three items is processed by reading up to the first character that is not white space (which remains unread), or until no more characters can be read.

*Tm* Specifies the structure to contain the output of the **strptime** subroutine. If a conversion fails, the contents of the *Tm* structure are undefined.

## Return Values

If successful, the **strptime** subroutine returns a pointer to the character following the last character parsed. Otherwise, a null pointer is returned.

## Related Information

The **scanf** (“scanf, fscanf, sscanf, or wsscanf Subroutine” on page 109) subroutine, “strfmon Subroutine” on page 286, **strptime** (“strptime Subroutine” on page 288) subroutine, **time** subroutine, **wcsftime** (“wcsftime Subroutine” on page 447) subroutine.

LC\_TIME Category in the Locale Definition Source File Format in *AIX 5L Version 5.2 Files Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and List of Time and Monetary Formatting Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

## stty or gtty Subroutine

### Purpose

Sets or gets terminal state.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sgtty.h>
```

```
stty ( FileDescriptor, Buffer)
int FileDescriptor;
struct sgttyb *Buffer;
gtty (FileDescriptor, Buffer)
int FileDescriptor;
struct sgttyb *Buffer;
```

## Description

These subroutines have been made obsolete by the **ioctl** subroutine.

The **stty** subroutine sets the state of the terminal associated with the *FileDescriptor* parameter. The **gtty** subroutine retrieves the state of the terminal associated with *FileDescriptor*. To set the state of a terminal, the calling process must have write permission.

Use of the **stty** subroutine is equivalent to the **ioctl** (*FileDescriptor*, TIOSETP, *Buffer*) subroutine, while use of the **gtty** subroutine is equivalent to the **ioctl** (*FileDescriptor*, TIOGETP, *Buffer*) subroutine.

## Parameters

<i>FileDescriptor</i>	Specifies an open file descriptor.
<i>Buffer</i>	Specifies the buffer.

## Return Values

If the **stty** or **gtty** subroutine is successful, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Related Information

The **ioctl** subroutine.

The Input and Output Handling Programmer's Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## swab Subroutine

### Purpose

Copies bytes.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>

void swab ( From, To, NumberOfBytes)
const void *From;
void *To;
ssize_t NumberOfBytes;
```

## Description

The **swab** subroutine copies the number of bytes pointed to by the *NumberOfBytes* parameter from the location pointed to by the *From* parameter to the array pointed to by the *To* parameter, exchanging adjacent even and odd bytes.

The *NumberOfBytes* parameter should be even and nonnegative. If the *NumberOfBytes* parameter is odd and positive, the **swab** subroutine uses *NumberOfBytes* -1 instead. If the *NumberOfBytes* parameter is negative, the **swab** subroutine does nothing.

## Parameters

<i>From</i>	Points to the location of data to be copied.
<i>To</i>	Points to the array to which the data is to be copied.
<i>NumberOfBytes</i>	Specifies the number of even and nonnegative bytes to be copied.

## Related Information

The **memccpy**, **memchr**, **memcmp**, **memmove**, or **memset** subroutine, **string** (“strlen, strchr, strchr, strpbrk, strspn, strcspn, strstr, strtok, or strsep Subroutine” on page 291) subroutine.

List of Interfaces for Input and Output Handling in *AIX 5L Version 5.2 System Management Guide: Operating System and Devices*.

Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## swapoff Subroutine

### Purpose

Deactivates paging or swapping to a designated block device.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int swapoff (PathName)
char *PathName;
```

### Description

The **swapoff** subroutine deactivates a block device or logical volume that is actively being used for paging and swapping. There must be sufficient space to satisfy the system’s paging space requirements in the remaining devices after this device is deactivated or **swapoff** will fail.

### Parameters

*PathName* Specifies the full path name of the block device or logical volume.

### Error Codes

If an error occurs, the **errno** global variable is set to indicate the error:

**EBUSY** The deactivation is already running.

<b>EINTR</b>	The signal was received during the processing of a request.
<b>ENODEV</b>	The <i>PathName</i> file does not exist.
<b>ENOMEM</b>	No memory is available.
<b>ENOSPC</b>	There is not enough space in other paging spaces to satisfy the system's requirements.
<b>ENOTBLK</b>	The device must be a block device or logical volume.
<b>ENOTDIR</b>	A component of the <i>PathName</i> prefix is not a directory.
<b>EPERM</b>	Caller does not have proper authority.

Other errors are from calls to the device driver's **open** subroutine or **ioctl** subroutine.

## Related Information

The **swapoff** command.

The Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## swapon Subroutine

### Purpose

Activates paging or swapping to a designated block device.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int swapon ( PathName)
char *PathName;
```

### Description

The **swapon** subroutine makes the designated block device available to the system for allocation for paging and swapping.

The specified block device must be a logical volume on a disk device. The paging space size is determined from the current size of the logical volume.

### Parameters

*PathName* Specifies the full path name of the block device.

### Error Codes

If an error occurs, the **errno** global variable is set to indicate the error:

<b>EINTR</b>	Signal was received during processing of a request.
<b>EINVAL</b>	Invalid argument (size of device is invalid).
<b>ENOENT</b>	The <i>PathName</i> file does not exist.
<b>ENOMEM</b>	The maximum number of paging space devices (16) are already defined, or no memory is available.
<b>ENOTBLK</b>	Block device required.
<b>ENOTDIR</b>	A component of the <i>PathName</i> prefix is not a directory.

**ENXIO** No such device address.

Other errors are from calls to the device driver's **open** subroutine or **ioctl** subroutine.

## Related Information

The **swapoff** subroutine, **swapqry** subroutine.

The **swapoff** command, **swapon** command.

The Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## swapqry Subroutine

### Purpose

Returns paging device status.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int swapqry (PathName, Buffer)
char *PathName;
struct pginfo *Buffer;
```

### Description

The **swapqry** subroutine returns information to a user-designated buffer about active paging and swap devices.

### Parameters

*PathName* Specifies the full path name of the block device.  
*Buffer* Points to the buffer into which the status is stored.

### Return Values

The **swapqry** subroutine returns 0 if the *PathName* value is an active paging device. If the *Buffer* value is not null, it also returns status information.

### Error Codes

If an error occurs, the subroutine returns -1 and the **errno** global variable is set to indicate the error, as follows:

<b>EFAULT</b>	Buffer pointer is invalid.
<b>EINVAL</b>	Invalid argument.
<b>EINTR</b>	Signal was received while processing request.
<b>ENODEV</b>	Device is not an active paging device.
<b>ENOENT</b>	The <i>PathName</i> file does not exist.
<b>ENOTBLK</b>	Block device required.

**ENOTDIR** A component of the *PathName* prefix is not a directory.  
**ENXIO** No such device address.

## Related Information

The **swapoff** subroutine, **swapon** subroutine.

The **swapoff** command, **swapon** command.

Paging Space Overview in *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

Subroutines Overview and Understanding Paging Space Programming Requirements in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## symlink Subroutine

### Purpose

Makes a symbolic link to a file.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int symlink ( Path1, Path2)  
const char *Path1;  
const char *Path2;
```

### Description

The **symlink** subroutine creates a symbolic link with the file named by the *Path2* parameter, which refers to the file named by the *Path1* parameter.

As with a hard link (described in the **link** subroutine), a symbolic link allows a file to have multiple names. The presence of a hard link guarantees the existence of a file, even after the original name has been removed. A symbolic link provides no such assurance. In fact, the file named by the *Path1* parameter need not exist when the link is created. In addition, a symbolic link can cross file system boundaries.

When a component of a path name refers to a symbolic link rather than a directory, the path name contained in the symbolic link is resolved. If the path name in the symbolic link starts with a / (slash), it is resolved relative to the root directory of the process. If the path name in the symbolic link does not start with / (slash), it is resolved relative to the directory that contains the symbolic link.

If the symbolic link is not the last component of the original path name, remaining components of the original path name are resolved from the symbolic-link point.

If the last component of the path name supplied to a subroutine refers to a symbolic link, the symbolic link path name may or may not be traversed. Most subroutines always traverse the link; for example, the **chmod**, **chown**, **link**, and **open** subroutines. The **statx** subroutine takes an argument that determines whether the link is to be traversed.

The following subroutines refer only to the symbolic link itself, rather than to the object to which the link refers:

<b>mkdir</b>	Fails with the <b>EEXIST</b> error code if the target is a symbolic link.
<b>mknod</b>	Fails with the <b>EEXIST</b> error code if a symbolic link exists with the same name as the target file as specified by the <i>Path</i> parameter in the <b>mknod</b> and <b>mkfifo</b> subroutines.
<b>open</b>	Fails with <b>EEXIST</b> error code when the <b>O_CREAT</b> and <b>O_EXCL</b> flags are specified and a symbolic link exists for the path name specified.
<b>readlink</b> (“readlink Subroutine” on page 22)	Applies only to symbolic links.
<b>rename</b> (“rename Subroutine” on page 42)	Renames the symbolic link if the file to be renamed (the <i>FromPath</i> parameter for the <b>rename</b> subroutine) is a symbolic link. If the new name (the <i>ToPath</i> parameter for the <b>rename</b> subroutine) refers to an existing symbolic link, the symbolic link is destroyed.
<b>rmdir</b> (“rmdir Subroutine” on page 48)	Fails with the <b>ENOTDIR</b> error code if the target is a symbolic link.
<b>symlink</b>	Running this subroutine causes an error if a symbolic link named by the <i>Path2</i> parameter already exists. A symbolic link can be created that refers to another symbolic link; that is, the <i>Path1</i> parameter can refer to a symbolic link.
<b>unlink</b> (“unlink Subroutine” on page 423)	Removes the symbolic link.

Since the mode of a symbolic link cannot be changed, its mode is ignored during the lookup process. Any files and directories referenced by a symbolic link are checked for access normally.

## Parameters

<i>Path1</i>	Specifies the contents of the <i>Path2</i> symbolic link. This value is a null-terminated string representing the object to which the symbolic link will point. <i>Path1</i> cannot be the null value and cannot be more than <b>PATH_MAX</b> characters long. <b>PATH_MAX</b> is defined in the <b>limits.h</b> file.
<i>Path2</i>	Names the symbolic link to be created.

## Return Values

Upon successful completion, the **symlink** subroutine returns a value of 0. If the **symlink** subroutine fails, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **symlink** subroutine fails if one or more of the following are true:

<b>EEXIST</b>	<i>Path2</i> already exists.
<b>EACCES</b>	The requested operation requires writing in a directory with a mode that denies write permission.
<b>EROFS</b>	The requested operation requires writing in a directory on a read-only file system.
<b>ENOSPC</b>	The directory in which the entry for the symbolic link is being placed cannot be extended because there is no space left on the file system containing the directory.
<b>EDQUOT</b>	The directory in which the entry for the new symbolic link is being placed cannot be extended or disk blocks could not be allocated for the symbolic link because the user’s or group’s quota of disk blocks on the file system containing the directory has been exhausted.

The **symlink** subroutine can be unsuccessful for other reasons. See "Base Operating System Error Codes For Services That Require Path-Name Resolution" for a list of additional errors.

## Related Information

The **chown**, **fchown**, **chownx**, or **fchown** subroutine, **link** subroutine, **mkdir** subroutine, **mknod** subroutine, **openx**, **open**, or **create** subroutine, **readlink** ("readlink Subroutine" on page 22) subroutine, **rename** ("rename Subroutine" on page 42) subroutine, **rmdir** ("rmdir Subroutine" on page 48) subroutine, **statx** ("statx, stat, lstat, fstatx, fstat, fullstat, fullstat, stat64, lstat64, or fstat64 Subroutine" on page 277) subroutine, **unlink** ("unlink Subroutine" on page 423) subroutine.

The **ln** command.

The **limits.h** file.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## sync Subroutine

### Purpose

Updates all file systems.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
void sync ( )
```

### Description

The **sync** subroutine causes all information in memory that should be on disk to be written out. The writing, although scheduled, is not necessarily complete upon return from this subroutine. Types of information to be written include modified superblocks, i-nodes, data blocks, and indirect blocks.

The **sync** subroutine should be used by programs that examine a file system, such as the **df** and **fsck** commands.

If Network File System (NFS) is installed on your system, information in memory that relates to remote files is scheduled to be sent to the remote node.

## Related Information

The **fsync** subroutine.

The **df** command, **sync** command.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## **`_sync_cache_range` Subroutine**

### **Purpose**

Synchronizes the I cache with the D cache.

### **Library**

Standard C Library (**libc.a**)

### **Syntax**

```
void _sync_cache_range (eaddr, count)
caddr_t eaddr;
uint count;
```

### **Description**

The **`_sync_cache_range`** subroutine synchronizes the I cache with the D cache, given an effective address and byte count. Programs performing instruction modification can call this routine to ensure that the most recent instructions are fetched for the address range.

### **Parameters**

*eaddr* Specifies the starting effective address of the address range.  
*count* Specifies the byte count of the address range.

### **Related Information**

The **`clf`** (Cache Line Flush) Instruction in *Assembler Language Reference*.

---

## **`sysconf` Subroutine**

### **Purpose**

Determines the current value of a specified system limit or option.

### **Library**

Standard C Library (**libc.a**)

### **Syntax**

```
#include <unistd.h>

long int sysconf ( Name)
int Name;
```

### **Description**

The **`sysconf`** subroutine determines the current value of certain system parameters, the configurable system limits, or whether optional features are supported. The *Name* parameter represents the system variable to be queried.

## Parameters

Name	Specifies which system variable setting should be returned. The valid values for the <i>Name</i> parameter are defined in the <b>limits.h</b> , <b>time.h</b> , and <b>unistd.h</b> files and are described below:
<code>_SC_AIO_LISTIO_MAX</code>	Maximum number of Input and Output operations that can be specified in a list Input and Output call.
<code>_SC_AIO_MAX</code>	Maximum number of outstanding asynchronous Input and Output operations.
<code>_SC_ASYNCHRONOUS_IO</code>	Implementation supports the Asynchronous Input and Output option.
<code>_SC_ARG_MAX</code>	Specifies the maximum byte length of the arguments for one of the <b>exec</b> functions, including environment data.
<code>_SC_BC_BASE_MAX</code>	Specifies the maximum number <b>ibase</b> and <b>obase</b> variables allowed by the <b>bc</b> command.
<code>_SC_BC_DIM_MAX</code>	Specifies the maximum number of elements permitted in an array by the <b>bc</b> command.
<code>_SC_BC_SCALE_MAX</code>	Specifies the maximum <b>scale</b> variable allowed by the <b>bc</b> command.
<code>_SC_BC_STRING_MAX</code>	Specifies the maximum length of a string constant allowed by the <b>bc</b> command.
<code>_SC_CHILD_MAX</code>	Specifies the number of simultaneous processes per real user ID.
<code>_SC_CLK_TCK</code>	Indicates the clock-tick increment as defined by the <b>CLK_TCK</b> in the <b>time.h</b> file.
<code>_SC_COLL_WEIGHTS_MAX</code>	Specifies the maximum number of weights that can be assigned to an entry of the <b>LC_COLLATE</b> keyword in the locale definition file.
<code>_SC_DELAYTIMER_MAX</code>	Maximum number of Timer expiration overruns.
<code>_SC_EXPR_NEST_MAX</code>	Specifies the maximum number of expressions that can be nested within parentheses by the <b>expr</b> command.
<code>_SC_JOB_CONTROL</code>	If this symbol is defined, job control is supported.
<code>_SC_IOV_MAX</code>	Specifies the maximum number of iovec structures one process has available for use with the <b>readv</b> and <b>writv</b> subroutines.
<code>_SC_LARGE_PAGESIZE</code>	Size (in bytes) of a large-page.
<code>_SC_LINE_MAX</code>	Specifies the maximum byte length of a command's input line (either standard input or another file) when a command is described as processing text files. The length includes room for the trailing new-line character.
<code>_SC_LOGIN_NAME_MAX</code>	Maximum length of a login name.
<code>_SC_MQ_OPEN_MAX</code>	Maximum number of open message queue descriptors.
<code>_SC_MQ_PRIO_MAX</code>	Maximum number of message priorities.
<code>_SC_MEMLOCK</code>	Implementation supports the Process Memory Locking option.
<code>_SC_MEMLOCK_RANGE</code>	Implementation supports the Range Memory Locking option.
<code>_SC_MEMORY_PROTECTION</code>	Implementation supports the Memory Protection option.
<code>_SC_MESSAGE_PASSING</code>	Implementation supports the Message Passing option.
<code>_SC_NGROUPS_MAX</code>	Specifies the maximum number of simultaneous supplementary group IDs per process.
<code>_SC_OPEN_MAX</code>	Specifies the maximum number of files that one process can have open at any one time.
<code>_SC_PASS_MAX</code>	Specifies the maximum number of significant characters in a password (not including the terminating null character).
<code>_SC_PASS_MAX</code>	Maximum number of significant bytes in a password.
<code>_SC_PAGESIZE</code>	Equivalent to <code>_SC_PAGE_SIZE</code> .
<code>_SC_PAGE_SIZE</code>	Size in bytes of a page.
<code>_SC_PRIORITIZED_IO</code>	Implementation supports the Prioritized Input and Output option.
<code>_SC_PRIORITY_SCHEDULING</code>	Implementation supports the Process Scheduling option.
<code>_SC_RE_DUP_MAX</code>	Specifies the maximum number of repeated occurrences of a regular expression permitted when using the $\{ m, n \}$ interval notation.
<code>_SC_RTSIG_MAX</code>	Maximum number of Realtime Signals reserved for applications use.
<code>_SC_REALTIME_SIGNALS</code>	Implementation supports the Realtime Signals Extension option.

<code>_SC_SAVED_IDS</code>	If this symbol is defined, each process has a saved set-user ID and set-group ID.
<code>_SC_SEM_NSEMS_MAX</code>	Maximum number of Semaphores per process.
<code>_SC_SEM_VALUE_MAX</code>	Maximum value a Semaphore may have.
<code>_SC_SEMAPHORES</code>	Implementation supports the Semaphores option.
<code>_SC_SHARED_MEMORY_OBJECTS</code>	Implementation supports the Shared Memory Objects option.
<code>_SC_SIGQUEUE_MAX</code>	Maximum number of signals a process may send and have pending at any time.
<code>_SC_STREAM_MAX</code>	Specifies the maximum number of streams that one process can have open simultaneously.
<code>_SC_SYNCHRONIZED_IO</code>	Implementation supports the Synchronised Input and Output option.
<code>_SC_TIMER_MAX</code>	Maximum number of per-process Timers.
<code>_SC_TIMERS</code>	Implementation supports the Timers option.
<code>_SC_TZNAME_MAX</code>	Specifies the maximum number of bytes supported for the name of a time zone (not of the <b>TZ</b> value).
<code>_SC_VERSION</code>	Indicates that the version or revision number of the POSIX standard is implemented to indicate the 4-digit year and 2-digit month that the standard was approved by the IEEE Standards Board. This value is currently the long integer 198808.
<code>_SC_XBS5_ILP32_OFF32</code>	Implementation provides a C-language compilation environment with 32-bit int, long, pointer and <code>off_t</code> types.
<code>_SC_XBS5_ILP32_OFFBIG</code>	Implementation provides a C-language compilation environment with 32-bit int, long and pointer types and an <code>off_t</code> type using at least 64 bits.
<code>_SC_XBS5_LP64_OFF64</code>	Implementation provides a C-language compilation environment with 32-bit int and 64-bit long, pointer and <code>off_t</code> types.
<code>_SC_XBS5_LPBIG_OFFBIG</code>	Implementation provides a C-language compilation environment with an int type using at least 32 bits and long, pointer and <code>off_t</code> types using at least 64 bits.
<code>_SC_XOPEN_CRYPT</code>	Indicates that the system supports the X/Open Encryption Feature Group.
<code>_SC_XOPEN_LEGACY</code>	The implementation supports the Legacy Feature Group.
<code>_SC_XOPEN_REALTIME</code>	The implementation supports the X/Open Realtime Feature Group.
<code>_SC_XOPEN_REALTIME_THREADS</code>	The implementation supports the X/Open Realtime Threads Feature Group.
<code>_SC_XOPEN_ENH_I18N</code>	Indicates that the system supports the X/Open Enhanced Internationalization Feature Group.
<code>_SC_XOPEN_SHM</code>	Indicates that the system supports the X/Open Shared Memory Feature Group.
<code>_SC_XOPEN_VERSION</code>	Indicates that the version or revision number of the X/Open standard is implemented.
<code>_SC_XOPEN_XCU_VERSION</code>	Specifies the value describing the current version of the XCU specification.
<code>_SC_ATEXIT_MAX</code>	Specifies the maximum number of register functions for the <code>atexit</code> subroutine.
<code>_SC_PAGE_SIZE</code>	Specifies page-size granularity of memory.
<code>_SC_AES_OS_VERSION</code>	Indicates OSF AES version.
<code>_SC_2_VERSION</code>	Specifies the value describing the current version of POSIX.2.
<code>_SC_2_C_BIND</code>	Indicates that the system supports the C Language binding option.
<code>_SC_2_C_CHAR_TERM</code>	Indicates that the system supports at least one terminal type.
<code>_SC_2_C_DEV</code>	Indicates that the system supports the C Language Development Utilities Option.
<code>_SC_2_C_VERSION</code>	Specifies the value describing the current version of POSIX.2 with the C Language binding.
<code>_SC_2_FORT_DEV</code>	Indicates that the system supports the FORTRAN Development Utilities Option.
<code>_SC_2_FORT_RUN</code>	Indicates that the system supports the FORTRAN Development Utilities Option.
<code>_SC_2_LOCALEDEF</code>	Indicates that the system supports the creation of locales.
<code>_SC_2_SW_DEV</code>	Indicates that the system supports the Software Development Utilities Option.
<code>_SC_2_UPE</code>	Indicates that the system supports the User Portability Utilities Option.
<code>_SC_NPROCESSORS_CONF</code>	Number of processors configured.
<code>_SC_NPROCESSORS_ONLN</code>	Number of processors online.
<code>_SC_THREAD_DATAKEYS_MAX</code>	Maximum number of data keys that can be defined in a process.

<code>_SC_THREAD_DESTRUCTOR_ITERATIONS</code>	Maximum number of attempts made to destroy a thread's thread-specific data.
<code>_SC_THREAD_KEYS_MAX</code>	Maximum number of data keys per process.
<code>_SC_THREAD_STACK_MIN</code>	Minimum value for the threads stack size.
<code>_SC_THREAD_THREADS_MAX</code>	Maximum number of threads within a process.
<code>_SC_REENTRANT_FUNCTIONS</code>	System supports reentrant functions (reentrant functions must be used in multi-threaded applications).
<code>_SC_THREADS</code>	System supports POSIX threads.
<code>_SC_THREAD_ATTR_STACKADDR</code>	System supports the stack address option for POSIX threads (stackaddr attribute of threads).
<code>_SC_THREAD_ATTR_STACKSIZE</code>	System supports the stack size option for POSIX threads (stacksize attribute of threads).
<code>_SC_THREAD_PRIORITY_SCHEDULING</code>	System supports the priority scheduling for POSIX threads.
<code>_SC_THREAD_PRIO_INHERIT</code>	System supports the priority inheritance protocol for POSIX threads (priority inversion protocol for mutexes).
<code>_SC_THREAD_PRIO_PROTECT</code>	System supports the priority ceiling protocol for POSIX threads (priority inversion protocol for mutexes).
<code>_SC_THREAD_PROCESS_SHARED</code>	System supports the process sharing option for POSIX threads (pshared attribute of mutexes and conditions).
<code>_SC_TTY_NAME_MAX</code>	Maximum length of a terminal device name.
<b>Note:</b> The <code>_SYNCHRONIZED_IO</code> , <code>_SC_FSYNC</code> , and <code>SC_MAPPED_FILES</code> commands apply to operating system version 4.3 and later releases.	
<code>_SC_SYNCHRONIZED_IO</code>	Implementation supports the <b>Synchronized Input and Output</b> option.
<code>_SC_FSYNC</code>	Implementation supports the <b>File Synchronization</b> option.
<code>_SC_MAPPED_FILES</code>	Implementation supports the <b>Memory Mapped Files</b> option.
<code>_SC_LPAR_ENABLED</code>	Indicates whether LPARs are enabled or not.
<code>_SC_AIX_KERNEL_BITMODE</code>	Determines if the kernel is 32-bit or 64-bit.
<code>_SC_AIX_REALMEM</code>	Determines the amount of real memory in kilobytes.
<code>_SC_AIX_HARDWARE_BITMODE</code>	Determines whether the machine is 32-bit or 64-bit.
<code>_SC_AIX_MP_CAPABLE</code>	Determines if the hardware is MP-capable or not.
<b>Note:</b> The <code>_SC_AIX_MP_CAPABLE</code> variable is available only to the root user.	

The values returned for the variables supported by the system do not change during the lifetime of the process making the call.

## Return Values

If the `sysconf` subroutine is successful, the current value of the system variable is returned. The returned value cannot be more restrictive than the corresponding value described to the application by the `limits.h`, `time.h`, or `unistd.h` file at compile time. The returned value does not change during the lifetime of the calling process. If the `sysconf` subroutine is unsuccessful, a value of -1 is returned.

## Error Codes

If the `Name` parameter is invalid, a value of -1 is returned and the `errno` global variable is set to indicate the error. If the `Name` parameter is valid but is a variable not supported by the system, a value of -1 is returned, and the `errno` global variable is set to a value of `EINVAL`. If the system variable `_SC_AIX_MP_CAPABLE` is accessed by a non-root user, a value of -1 is returned and the `errno` global variable indicates the error.

## File

`/usr/include/limits.h`

Contains system-defined limits.

## Related Information

The **confstr** subroutine, **pathconf** subroutine.

The **bc** command, **expr** command.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## sysconfig Subroutine

### Purpose

Provides a service for controlling system/kernel configuration.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int sysconfig ( Cmd, Parmp, ParmLen)
int Cmd;
void *Parmp;
int ParmLen;
```

### Description

The **sysconfig** subroutine is used to customize the operating system. This subroutine provides a means of loading, unloading, and configuring kernel extensions. These kernel extensions can be additional kernel services, system calls, device drivers, or file systems. The **sysconfig** subroutine also provides the ability to read and set system run-time operating parameters.

Use of the **sysconfig** subroutine requires appropriate privilege.

The particular operation that the **sysconfig** subroutine provides is defined by the value of the *Cmd* parameter. The following operations are defined:

<b>SYS_KLOAD</b> ("SYS_KLOAD sysconfig Operation" on page 321)	Loads a kernel extension object file into kernel memory.
<b>SYS_SINGLELOAD</b> ("SYS_SINGLELOAD sysconfig Operation" on page 328)	Loads a kernel extension object file only if it is not already loaded.
<b>SYS_QUERYLOAD</b> ("SYS_QUERYLOAD sysconfig Operation" on page 326)	Determines if a specified kernel object file is loaded.
<b>SYS_KULOAD</b> ("SYS_KULOAD sysconfig Operation" on page 323)	Unloads a previously loaded kernel object file.
<b>SYS_QDVSW</b> ("SYS_QDVSW sysconfig Operation" on page 325)	Checks the status of a device switch entry in the device switch table.

<b>SYS_CFGDD</b> ("SYS_CFGDD sysconfig Operation")	Calls the specified device driver configuration routine (module entry point).
<b>SYS_CFGKMOD</b> ("SYS_CFGKMOD sysconfig Operation" on page 318)	Calls the specified module at its module entry point for configuration purposes.
<b>SYS_GETPARMS</b> ("SYS_GETPARMS sysconfig Operation" on page 320)	Returns a structure containing the current values of run-time system parameters found in the <b>var</b> structure.
<b>SYS_SETPARMS</b> ("SYS_SETPARMS sysconfig Operation" on page 327)	Sets run-time system parameters from a caller-provided structure.
<b>SYS_GETLPARINFO</b> ("SYS_GETLPAR_INFO sysconfig Operation" on page 319)	Copies the system LPAR information into a user-allocated buffer.

In addition, the **SYS\_64BIT** flag can be bitwise or'ed with the *Cmd* parameter (if the *Cmd* parameter is **SYS\_KLOAD** or **SYS\_SINGLELOAD**). For kernel extensions, this indicates that the kernel extension does not export 64-bit system calls, but that all 32-bit system calls also work for 64-bit applications. For device drivers, this indicates that the device driver can be used by 64-bit applications.

"Loader Symbol Binding Support" on page 322 explains the symbol binding support provided when loading kernel object files.

## Parameters

<i>Cmd</i>	Specifies the function that the <b>sysconfig</b> subroutine is to perform.
<i>Parmp</i>	Specifies a user-provided structure.
<i>Parmlen</i>	Specifies the length of the user-provided structure indicated by the <i>Parmp</i> parameter.

## Return Values

These sysconfig operations return a value of 0 upon successful completion of the subroutine. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

Any sysconfig operation requiring a structure from the caller fails if the structure is not entirely within memory addressable by the calling process. A return value of -1 is passed back and the **errno** global variable is set to **EFAULT**.

## Related Information

The **ddconfig** device driver entry point.

Device Configuration Subsystem, Kernel Environment, Understanding Kernel Extension Binding in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

---

## SYS\_CFGDD sysconfig Operation

### Purpose

Calls a previously loaded device driver at its module entry point.

## Description

The **SYS\_CFGDD** sysconfig operation calls a previously loaded device driver at its module entry point. The device driver's module entry point, by convention, is its **ddconfig** entry point. The **SYS\_CFGDD** operation is typically invoked by device configure or unconfigure methods to initialize or terminate a device driver, or to request device vital product data.

The **sysconfig** subroutine puts no restrictions on the command code passed to the device driver. This allows the device driver's **ddconfig** entry point to provide additional services, if desired.

The *parmp* parameter on the **SYS\_CFGDD** operation points to a **cfg\_dd** structure defined in the **sys/sysconfig.h** file. The *parmlen* parameter on the **sysconfig** system call should be set to the size of this structure.

If the *kmid* variable in the **cfg\_dd** structure is 0, the desired device driver is assumed to be already installed in the device switch table. The major portion of the device number (passed in the *devno* field in the **cfg\_dd** structure) is used as an index into the device switch table. The device switch table entry indexed by this *devno* field contains the device driver's **ddconfig** entry point to be called.

If the *kmid* variable is not 0, it contains the module ID to use in calling the device driver. A **uio** structure is used to pass the address and length of the device-dependent structure, specified by the *cfg\_dd.ddsptr* and *cfg\_dd.ddslen* fields, to the device driver being called.

The **ddconfig** device driver entry point provides information on how to define the **ddconfig** subroutine.

The device driver to be called is responsible for using the appropriate routines to copy the device-dependent structure (DDS) from user to kernel space.

## Return Values

If the **SYS\_CFGDD** operation successfully calls the specified device driver, the return code from the **ddconfig** subroutine determines the value returned by this subroutine. If the **ddconfig** routine's return code is 0, then the value returned by the **sysconfig** subroutine is 0. Otherwise the value returned is a -1, and the **errno** global variable is set to the return code provided by the device driver **ddconfig** subroutine.

## Error Codes

Errors detected by the **SYS\_CFGDD** operation result in the following values for the **errno** global variable:

<b>EACCESS</b>	The calling process does not have the required privilege.
<b>EFAULT</b>	The calling process does not have sufficient authority to access the data area described by the <i>parmp</i> and <i>parmlen</i> parameters provided on the system call. This error is also returned if an I/O error occurred when accessing data in this area.
<b>EINVAL</b>	Invalid module ID.
<b>ENODEV</b>	Module ID specified by the <b>cfg_dd.kmid</b> field was 0, and an invalid or undefined <b>devno</b> value was specified.

## Related Information

The **sysconfig** ("sysconfig Subroutine" on page 315) subroutine.

The **ddconfig** device driver entry point.

The **uio** structure.

Device Configuration Subsystem Programming Introduction, Device Dependent Structure (DDS) Overview, Device Driver Kernel Extension Overview, Programming in the Kernel Environment Overview,

## SYS\_CFGKMOD sysconfig Operation

### Purpose

Invokes a previously loaded kernel object file at its module entry point.

### Description

The **SYS\_CFGKMOD** sysconfig operation invokes a previously loaded kernel object file at its module entry point, typically for initialization or termination functions. The **SYS\_CFGDD** (“SYS\_CFGDD sysconfig Operation” on page 316) operation performs a similar function for device drivers.

The *parm* parameter on the **sysconfig** subroutine points to a **cfg\_kmod** structure, which is defined in the **sys/sysconfig.h** file. The *kmid* field in this structure specifies the kernel module ID of the module to invoke. This value is returned when using the **SYS\_KLOAD** (“SYS\_KLOAD sysconfig Operation” on page 321) or **SYS\_SINGLELOAD** (“SYS\_SINGLELOAD sysconfig Operation” on page 328) operation to load the object file.

The *cmd* field in the **cfg\_kmod** structure is a module-dependent parameter specifying the action that the routine at the module’s entry point should perform. This is typically used for initialization and termination commands after loading and prior to unloading the object file.

The *mdiptr* field in the **cfg\_kmod** structure points to a module-dependent structure whose size is specified by the *mdilen* field. This field is used to provide module-dependent information to the module to be called. If no such information is needed, the *mdiptr* field can be null.

If the *mdiptr* field is not null, then the **SYS\_CFGKMOD** operation builds a **uio** structure describing the address and length of the module-dependent information in the caller’s address space. The *mdiptr* and *mdilen* fields are used to fill in the fields of this **uio** structure. The module is then called at its module entry point with the *cmd* parameter and a pointer to the **uio** structure. If there is no module-dependent information to be provided, the *uiop* parameter passed to the module’s entry point is set to null.

The module’s entry point should be defined as follows:

```
int module_entry(cmd, uiop)
int cmd;
struct uio *uiop;
```

The definition of the module-dependent information and its length is specific to the module being configured. The called module is responsible for using the appropriate routines to copy the module-dependent information from user to kernel space.

### Return Values

If the kernel module to be invoked is successfully called, its return code determines the value that is returned by the **SYS\_CFGKMOD** operation. If the called module’s return code is 0, then the value returned by the **sysconfig** subroutine is 0. Otherwise the value returned is -1 and the **errno** global variable is set to the called module’s return code.

### Error Codes

Errors detected by the **SYS\_CFGKMOD** operation result in the following values for the **errno** global variable:

**EINVAL** Invalid module ID.

**EACCESS** The calling process does not have the required privilege.  
**EFAULT** The calling process does not have sufficient authority to access the data area described by the *parmp* and *parmlen* parameters provided on the system call. This error is also returned if an I/O error occurred when accessing data in this area.

## File

**sys/sysconfig.h** Contains structure definitions.

## Related Information

The **sysconfig** (“sysconfig Subroutine” on page 315) subroutine.

The **SYS\_CFGDD** (“SYS\_CFGDD sysconfig Operation” on page 316) sysconfig operation, **SYS\_KLOAD** (“SYS\_KLOAD sysconfig Operation” on page 321) sysconfig operation, **SYS\_SINGLELOAD** (“SYS\_SINGLELOAD sysconfig Operation” on page 328) sysconfig operation.

The **uio** structure.

Device Configuration Subsystem Programming Introduction in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*

Programming in the Kernel Environment Overview in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*

Understanding Kernel Extension Binding in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*

---

## SYS\_GETLPAR\_INFO sysconfig Operation

### Purpose

Copies the system LPAR information into a user-allocated buffer.

### Description

The **SYS\_GETLPAR\_INFO** sysconfig operation copies the system LPAR information into a user-allocated buffer.

The *parmp* parameter on the **sysconfig** subroutine points to a structure of type **getlpar\_info**. Within the **getlpar\_info** structure, the *lpar\_namelen* field must be set by the user to the maximum length of the character buffer pointed to by *lpar\_name*. On return, the *lpar\_namelen* field will have its value replaced by the actual length of the *lpar\_name* field. However, only the minimum of the actual length or the length provided by the user will be copied into the buffer pointed to by *lpar\_name*. The *lpar\_namesz*, *lpar\_num*, and *lpar\_name* fields will contain valid data on returning from the call only if the system is running as an LPAR as indicated by the value of the *lpar\_flags* field being equal to **LPAR\_ENABLED**.

If a value of 0 is specified for the *lpar\_namesz* field, the partition name will not be copied out.

If the system is not an LPAR (namely it is running as an SMP system), but it is LPAR-capable, the **LPAR\_CAPABLE** flag will be set on return.

The **getlpar\_info** structure is defined below:

<i>lpar_flags</i>	unsigned short	<b>LPAR_ENABLED:</b> System is LPAR enabled. <b>LPAR_CAPABLE:</b> System is LPAR capable, but running in SMP mode.
<i>lpar_namesz</i>	unsigned short	Size of partition name.
<i>lpar_num</i>	int	Partition Number.
<i>lpar_name</i>	char *	Partition Name.

**Note:** The *parmlen* parameter (which is the third parameter to the **sysconfig** system call) is ignored by the **SYS\_GETLPAR\_INFO** sysconfig operation.

## Error Codes

The **SYS\_GETLPAR\_INFO** operation returns a value of -1 if an error occurs and the **errno** global variable is set to one of the following error codes:

<b>EFAULT</b>	The calling process does not have sufficient authority to access the data area described by the <i>parmp</i> and <i>parmlen</i> parameters provided on the subroutine or the <i>lpar_name</i> field in the <b>getlpar_info</b> structure. This error is also returned if an I/O error occurred when accessing data in any of these areas.
<b>EINVAL</b>	Invalid command parameter to the <b>sysconfig</b> subroutine.

## Files

**sys/sysconfig.h** Contains structure definitions and flags.

## Related Information

The “sysconfig Subroutine” on page 315.

Programming in the Kernel Environment Overview in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

---

## SYS\_GETPARMS sysconfig Operation

### Purpose

Copies the system parameter structure into a user-specified buffer.

### Description

The **SYS\_GETPARMS** sysconfig operation copies the system parameter **var** structure into a user-allocated buffer. This structure may be used for informational purposes alone or prior to setting specific system parameters.

In order to set system parameters, the required fields in the **var** structure must be modified, and then the **SYS\_SETPARMS** (“SYS\_SETPARMS sysconfig Operation” on page 327) operation can be called to change the system run-time operating parameters to the desired state.

The *parmp* parameter on the **sysconfig** subroutine points to a buffer that is to contain all or part of the **var** structure defined in the **sys/var.h** file. The fields in the **var\_hdr** part of the **var** structure are used for parameter update control.

The *parmlen* parameter on the system call should be set to the length of the **var** structure or to the number of bytes of the structure that is desired. The complete definition of the system parameters structure can be found in the **sys/var.h** file.

## Return Values

The **SYS\_GETPARMS** operation returns a value of -1 if an error occurs and the **errno** global variable is set to one of the following error codes.

## Error Codes

- EACCES** The calling process does not have the required privilege.
- EFAULT** The calling process does not have sufficient authority to access the data area described by the *parm* and *parmlen* parameters provided on the subroutine. This error is also returned if an I/O error occurred when accessing data in this area.

## File

**sys/var.h** Contains structure definitions.

## Related Information

The **sysconfig** (“sysconfig Subroutine” on page 315) subroutine and **sys\_parm** (“sys\_parm Subroutine” on page 336) subroutine.

The **SYS\_SETPARMS** (“SYS\_SETPARMS sysconfig Operation” on page 327) sysconfig operation.

Programming in the Kernel Environment Overview in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

---

## SYS\_KLOAD sysconfig Operation

### Purpose

Loads a kernel extension into the kernel.

### Description

The **SYS\_KLOAD** sysconfig operation is used to load a kernel extension object file specified by a path name into the kernel. A kernel module ID for that instance of the module is returned. The **SYS\_KLOAD** operation loads a new copy of the object file into the kernel even though one or more copies of the specified object file may have already been loaded into the kernel. The returned module ID can then be used for any of these three functions:

- Subsequent invocation of the module’s entry point (using the **SYS\_CFGKMOD** (“SYS\_CFGKMOD sysconfig Operation” on page 318) operation)
- Invocation of a device driver’s **ddconfig** subroutine (using the **SYS\_CFGDD** (“SYS\_CFGDD sysconfig Operation” on page 316) operation)
- Unloading the kernel module (using the **SYS\_KULOAD** (“SYS\_KULOAD sysconfig Operation” on page 323) operation).

The *parm* parameter on the **sysconfig** subroutine must point to a **cfg\_load** structure, (defined in the **sys/sysconfig.h** file), with the path field specifying the path name for a valid kernel object file. The *parmlen* parameter should be set to the size of the **cfg\_load** structure.

**Note:** A separate **sysconfig** operation, the **SYS\_SINGLELOAD** (“SYS\_SINGLELOAD sysconfig Operation” on page 328) operation, also loads kernel extensions. This operation, however, only loads the requested object file if not already loaded.

## Loader Symbol Binding Support

The following information describes the symbol binding support provided when loading kernel object files.

### Importing Symbols

Symbols imported from the kernel name space are resolved with symbols that exist in the kernel name space at the time of the load. (Symbols are imported from the kernel name space by specifying the `#!/unix` character string as the first field in an import list at link-edit time.)

Kernel modules can also import symbols from other kernel object files. These other kernel object files are loaded along with the specified object file if they are required to resolve the imported symbols.

**Finding Directory Locations for Unqualified File Names:** If the module header contains an unqualified base file name for the symbol (that is, no `/` [slash] characters in the name), a `libpath` search string is used to find the location of the shared object file required to resolve imported symbols. This `libpath` search string can be taken from one of two places. If the `libpath` field in the `cfg_load` structure is not null, then it points to a character string specifying the `libpath` to be used. However, if the `libpath` field is null, then the `libpath` is taken from the module header of the object file specified by the `path` field in the same (`cfg_load`) structure.

The `libpath` specification found in object files loaded in order to resolve imported symbols is not used.

The kernel loader service does not support deferred symbol resolution. The load of the kernel object file is terminated with an error if any imported symbols cannot be resolved.

### Exporting Symbols

Any symbols exported by the specified kernel object file are added to the kernel name space. This makes these symbols available to other subsequently loaded kernel object files. Any symbols specified with the **SYSCALL** keyword in the export list at link-edit time are added to the system call table at load time. These symbols are then available to application programs as a system call. Symbols can be added to the 32-bit and 64-bit system call tables separately by using the `syscall32` and `syscall64` keywords. Symbols can be added to both system call tables by using the `syscall3264` keyword. A kernel extension that just exports 32-bit system calls can have all its system calls exported to 64-bit as well by passing the `SYS_64BIT` flag or'ed with the `SYS_KLOAD` command to `sysconfig`.

Kernel object files loaded on behalf of the specified kernel object file to resolve imported symbols do not have their exported symbols added to the kernel name space.

These object files are considered private since they do not export symbols to the global kernel name space. For these types of object files, a new copy of the object file is loaded on each **SYS\_KLOAD** operation of a kernel extension that imports symbols from the private object file. In order for a kernel extension to add its exported symbols to the kernel name space, it must be explicitly loaded with the **SYS\_KLOAD** operation before any other object files using the symbols are loaded. For kernel extensions of this type (those exporting symbols to the kernel name space), typically only one copy of the object file should ever be loaded.

### Return Values

If the object file is loaded without error, the module ID is returned in the `kmid` variable within the `cfg_load` structure and the subroutine returns a value of **0**.

### Error Codes

On error, the subroutine returns a value of **-1** and the `errno` global variable is set to one of the following values:

<b>EACCESS</b>	One of the following reasons applies: <ul style="list-style-type: none"> <li>• The calling process does not have the required privilege.</li> <li>• An object module to be loaded is not an ordinary file.</li> <li>• The mode of the object module file denies read-only permission.</li> </ul>
<b>EFAULT</b>	The calling process does not have sufficient authority to access the data area described by the <i>parmp</i> and <i>parmlen</i> parameters provided on the system call. This error is also returned if an I/O error occurred when accessing data in this area.
<b>ENOEXEC</b>	The program file has the appropriate access permission, but has an invalid XCOFF object file indication in its header. The <b>SYS_KLOAD</b> operation only supports loading of XCOFF object files. This error is also returned if the loader is unable to resolve an imported symbol.
<b>EINVAL</b>	The program file has a valid XCOFF indicator in its header, but the header is damaged or is incorrect for the machine on which the file is to be run.
<b>ENOMEM</b>	The load requires more kernel memory than is allowed by the system-imposed maximum.
<b>ETXTBSY</b>	The object file is currently open for writing by some process.

## File

**sys/sysconfig.h**                      Contains structure definitions.

## Related Information

The **sysconfig** (“sysconfig Subroutine” on page 315) subroutine.

The **SYS\_SINGLELOAD** (“SYS\_SINGLELOAD sysconfig Operation” on page 328) sysconfig operation, **SYS\_KULOAD** (“SYS\_KULOAD sysconfig Operation”) sysconfig operation, **SYS\_CFGDD** (“SYS\_CFGDD sysconfig Operation” on page 316) sysconfig operation, **SYS\_CFGKMOD** (“SYS\_CFGKMOD sysconfig Operation” on page 318) sysconfig operation.

The **ddconfig** device driver entry point.

Device Configuration Subsystem Programming Introduction in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

Programming in the Kernel Environment Overview in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

Understanding Kernel Extension Binding in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

---

## SYS\_KULOAD sysconfig Operation

### Purpose

Unloads a loaded kernel object file and any imported kernel object files that were loaded with it.

### Description

The **SYS\_KULOAD** sysconfig operation unloads a previously loaded kernel file and any imported kernel object files that were automatically loaded with it. It does this by decrementing the load and use counts of the specified object file and any object file having symbols imported by the specified object file.

The *parmp* parameter on the **sysconfig** subroutine should point to a **cfg\_load** structure, as described for the **SYS\_KLOAD** (“SYS\_KLOAD sysconfig Operation” on page 321) operation. The *kmiid* field should specify the kernel module ID that was returned when the object file was loaded by the **SYS\_KLOAD** or

**SYS\_SINGLELOAD** (“SYS\_SINGLELOAD sysconfig Operation” on page 328) operation. The `path` and `libpath` fields are not used for this command and can be set to null. The `parmlen` parameter should be set to the size of the `cfg_load` structure.

Upon successful completion, the specified object file (and any other object files containing symbols that the specified object file imports) will have their load and use counts decremented. If there are no users of any of the module’s exports and its load count is 0, then the object file is immediately unloaded.

However, if there are users of this module (that is, modules bound to this module’s exported symbols), the specified module is not unloaded. Instead, it is unloaded on some subsequent unload request, when its use and load counts have gone to 0. The specified module is not in fact unloaded until all current users have been unloaded.

#### Notes:

1. Care must be taken to ensure that a subroutine has freed all of its system resources before being unloaded. For example, a device driver is typically prepared for unloading by using the **SYS\_CFGDD** (“SYS\_CFGDD sysconfig Operation” on page 316) operation and specifying termination.
2. If the use count is not 0, and you cannot force it to 0, the only way to terminate operation of the kernel extension is to reboot the machine.

“Loader Symbol Binding Support” on page 322 explains the symbol binding support provided when loading kernel object files.

## Return Values

If the unload operation is successful or the specified object file load count is successfully decremented, a value of 0 is returned.

## Error Codes

On error, the specified file and any imported files are not unloaded, nor are their load and use counts decremented. A value of -1 is returned and the `errno` global variable is set to one of the following:

<b>EACCESS</b>	The calling process does not have the required privilege.
<b>EINVAL</b>	Invalid module ID or the specified module is no longer loaded or already has a load count of 0.
<b>EFAULT</b>	The calling process does not have sufficient authority to access the data area described by the <code>parmp</code> and <code>parmlen</code> parameters provided to the subroutine. This error is also returned if an I/O error occurred when accessing data in this area.

## Related Information

The **SYS\_CFGDD** (“SYS\_CFGDD sysconfig Operation” on page 316) sysconfig operation, **SYS\_KLOAD** (“SYS\_KLOAD sysconfig Operation” on page 321) sysconfig operation, **SYS\_SINGLELOAD** (“SYS\_SINGLELOAD sysconfig Operation” on page 328) sysconfig operation.

The **sysconfig** (“sysconfig Subroutine” on page 315) subroutine.

Device Configuration Subsystem Programming Introduction in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

Programming in the Kernel Environment Overview in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

Understanding Kernel Extension Binding in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

---

## SYS\_QDVSW sysconfig Operation

### Purpose

Checks the status of a device switch entry in the device switch table.

### Description

The **SYS\_QDVSW** sysconfig operation checks the status of a device switch entry in the device switch table.

The *parm* parameter on the **sysconfig** subroutine points to a **qry\_devsw** structure defined in the **sys/sysconfig.h** file. The *parmlen* parameter on the subroutine should be set to the length of the **qry\_devsw** structure.

The *qry\_devsw* field in the **qry\_devsw** structure is modified to reflect the status of the device switch entry specified by the *qry\_devsw* field. ( Only the major portion of the *devno* field is relevant.) The following flags can be returned in the *status* field:

<b>DSW_UNDEFINED</b>	The device switch entry is not defined if this flag has a value of 0 on return.
<b>DSW_DEFINED</b>	The device switch entry is defined.
<b>DSW_CREAD</b>	The device driver in this device switch entry provides a routine for character reads or raw input. This flag is set when the device driver provides a <b>ddread</b> entry point.
<b>DSW_CWRITE</b>	The device driver in this device switch entry provides a routine for character writes or raw output. This flag is set when the device driver provides a <b>ddwrite</b> entry point.
<b>DSW_BLOCK</b>	The device switch entry is defined by a block device driver. This flag is set when the device driver provides a <b>ddstrategy</b> entry point.
<b>DSW_MPX</b>	The device switch entry is defined by a multiplexed device driver. This flag is set when the device driver provides a <b>ddmpx</b> entry point.
<b>DSW_SELECT</b>	The device driver in this device switch entry provides a routine for handling the <b>select</b> (“select Subroutine” on page 115) or <b>poll</b> subroutines. This flag is set when the device driver provides a <b>ddselect</b> entry point.
<b>DSW_DUMP</b>	The device driver defined by this device switch entry provides the capability to support one or more of its devices as targets for a kernel dump. This flag is set when the device driver has provided a <b>dddump</b> entry point.
<b>DSW_CONSOLE</b>	The device switch entry is defined by the console device driver.
<b>DSW_TCPATH</b>	The device driver in this device switch entry supports devices that are considered to be in the trusted computing path and provides support for the <b>revoke</b> (“revoke Subroutine” on page 45) and <b>frevoke</b> subroutines. This flag is set when the device driver provides a <b>ddrevoke</b> entry point.
<b>DSW_OPENED</b>	The device switch entry is defined and the device has outstanding opens. This flag is set when the device driver has at least one outstanding open.

The **DSW\_UNDEFINED** condition is indicated when the device switch entry has not been defined or has been defined and subsequently deleted. Multiple status flags may be set for other conditions of the device switch entry.

### Return Values

If no error is detected, this operation returns with a value of 0. If an error is detected, the return value is set to a value of -1.

### Error Codes

When an error is detected, the **errno** global variable is also set to one of the following values:

<b>EACCESS</b>	The calling process does not have the required privilege.
<b>EINVAL</b>	Device number exceeds the maximum allowed by the kernel.

**EFAULT** The calling process does not have sufficient authority to access the data area described by the *parmp* and *parmlen* parameters provided on the system call. This error is also returned if an I/O error occurred when accessing data in this area.

## File

**sys/sysconfig.h** Contains structure definitions.

## Related Information

The **sysconfig** (“sysconfig Subroutine” on page 315) subroutine.

The **ddread** device driver entry point, **ddwrite** device driver entry point, **ddstrategy** device driver entry point, **ddmpx** device driver entry point, **ddselect** device driver entry point, **dddump** device driver entry point, **ddrevoke** device driver entry point.

The **console** special file.

Device Configuration Subsystem Programming Introduction in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

Programming in the Kernel Environment Overview in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

Understanding Kernel Extension Binding in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

---

## SYS\_QUERYLOAD sysconfig Operation

### Purpose

Determines if a kernel object file has already been loaded.

### Description

The **SYS\_QUERYLOAD** sysconfig operation performs a query operation to determine if a given object file has been loaded. This object file is specified by the path field in the **cfg\_load** structure passed in with the *parmp* parameter. This operation utilizes the same **cfg\_load** structure that is specified for the **SYS\_KLOAD** (“SYS\_KLOAD sysconfig Operation” on page 321) operation.

If the specified object file is not loaded, the *kmid* field in the **cfg\_load** structure is set to a value of 0 on return. Otherwise, the kernel module ID of the module is returned in the *kmid* field. If multiple instances of the module have been loaded into the kernel, the module ID of the one most recently loaded is returned.

The *libpath* field in the **cfg\_load** structure is not used for this option.

**Note:** A path-name comparison is done to determine if the specified object file has been loaded. However, this operation will erroneously return a *not loaded* condition if the path name to the object file is expressed differently than it was on a previous load request.

“Loader Symbol Binding Support” on page 322 explains the symbol binding support provided when loading kernel object files.

## Return Values

If the specified object file is found, the module ID is returned in the *kmid* variable within the **cfg\_load** structure and the subroutine returns a 0. If the specified file is not found, a *kmid* variable of 0 is returned with a return code of 0.

## Error Codes

On error, the subroutine returns a -1 and the **errno** global variable is set to one of the following values:

- EACCES** The calling process does not have the required privilege.
- EFAULT** The calling process does not have sufficient authority to access the data area described by the *parm* and *parmlen* parameters provided on the subroutine. This error is also returned if an I/O error occurred when accessing data in this area.
- EFAULT** The *path* parameter points to a location outside of the allocated address space of the process.
- EIO** An I/O error occurred during the operation.

## Related Information

The **sysconfig** (“sysconfig Subroutine” on page 315) subroutine.

The **SYS\_SINGLELOAD** (“SYS\_SINGLELOAD sysconfig Operation” on page 328) sysconfig operation, **SYS\_KLOAD** (“SYS\_KLOAD sysconfig Operation” on page 321) sysconfig operation.

Programming in the Kernel Environment Overview in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

Understanding Kernel Extension Binding Overview in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

---

## SYS\_SETPARMS sysconfig Operation

### Purpose

Sets the kernel run-time tunable parameters.

### Description

The **SYS\_SETPARMS** sysconfig operation sets the current system parameters from a copy of the system parameter **var** structure provided by the caller. Only the run-time tunable parameters in the **var** structure can be set by this subroutine.

If the *var\_vers* and *var\_gen* values in the caller-provided structure do not match the *var\_vers* and *var\_gen* values in the current system **var** structure, no parameters are modified and an error is returned. The *var\_vers*, *var\_gen*, and *var\_size* fields in the structure should not be altered. The *var\_vers* value is assigned by the kernel and is used to insure that the correct version of the structure is being used. The **var\_gen** value is a generation number having a new value for each read of the structure. This provides consistency between the data read by the **SYS\_GETPARMS** (“SYS\_GETPARMS sysconfig Operation” on page 320) operation and the data written by the **SYS\_SETPARMS** operation.

The *parm* parameter on the **sysconfig** subroutine points to a buffer that contains all or part of the **var** structure as defined in the **sys/var.h** file.

The *parmlen* parameter on the subroutine should be set either to the length of the **var** structure or to the size of the structure containing the parameters to be modified. The number of system parameters modified by this operation is determined either by the *parmlen* parameter value or by the *var\_size* field in the caller-provided **var** structure. (The smaller of the two values is used.)

The structure provided by the caller must contain at least the header fields of the **var** structure. Otherwise, an error will be returned. Partial modification of a parameter in the **var** structure can occur if the caller's data area does not contain enough data to end on a field boundary. It is up to the caller to ensure that this does not happen.

## Return Values

The **SYS\_SETPARMS** sysconfig operation returns a value of -1 if an error occurred.

## Error Codes

When an error occurs, the **errno** global variable is set to one of the following values:

<b>EACCESS</b>	The calling process does not have the required privilege.
<b>EINVAL</b>	One of the following error situations exists: <ul style="list-style-type: none"><li>• The <code>var_vers</code> version number of the provided structure does not match the version number of the current <b>var</b> structure.</li><li>• The structure provided by the caller does not contain enough data to specify the header fields within the <b>var</b> structure.</li><li>• One of the specified variable values is invalid or not allowed. On the return from the subroutine, the <code>var_vers</code> field in the caller-provided buffer contains the byte offset of the first variable in the structure that was detected in error.</li></ul>
<b>EAGAIN</b>	The <code>var_gen</code> generation number in the structure provided does not match the current generation number in the kernel. This occurs if consistency is lost between reads and writes of this structure. The caller should repeat the read, modify, and write operations on the structure.
<b>EFAULT</b>	The calling process does not have sufficient authority to access the data area described by the <code>parmp</code> and <code>parmlen</code> parameters provided to the subroutine. This error is also returned if an I/O error occurred when accessing data in this area.

## File

**sys/var.h**            Contains structure definitions.

## Related Information

The **sysconfig** (“sysconfig Subroutine” on page 315) subroutine and **sys\_parm** (“sys\_parm Subroutine” on page 336) subroutine.

The **SYS\_GETPARMS** (“SYS\_GETPARMS sysconfig Operation” on page 320) sysconfig operation.

Programming in the Kernel Environment Overview in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

---

## SYS\_SINGLELOAD sysconfig Operation

### Purpose

Loads a kernel extension module if it is not already loaded.

### Description

The **SYS\_SINGLELOAD** sysconfig operation is identical to the **SYS\_KLOAD** (“SYS\_KLOAD sysconfig Operation” on page 321) operation, except that the **SYS\_SINGLELOAD** operation loads the object file only if an object file with the same path name has not already been loaded into the kernel.

If an object file with the same path name has already been loaded, the module ID for that object file is returned in the `kmid` field and its load count incremented. If the object file is not loaded, this operation performs the load request exactly as defined for the **SYS\_KLOAD** operation.

This option is useful in supporting global kernel routines where only one copy of the routine and its data can be present. Typically routines that export symbols to be added to the kernel name space are of this type.

**Note:** A path name comparison is done to determine if the same object file has already been loaded. However, this function will erroneously load a new copy of the object file into the kernel if the path name to the object file is expressed differently than it was on a previous load request.

“Loader Symbol Binding Support” on page 322 explains the symbol binding support provided when loading kernel object files.

## Return Values

The **SYS\_SINGLELOAD** operation returns the same set of error codes that the **SYS\_KLOAD** operation returns.

## Related Information

The **sysconfig** (“sysconfig Subroutine” on page 315) subroutine.

The **SYS\_KLOAD** (“SYS\_KLOAD sysconfig Operation” on page 321) sysconfig operation.

Programming in the Kernel Environment Overview, and Understanding Kernel Extension Binding in *AIX 5L Version 5.2 Kernel Extensions and Device Support Programming Concepts*.

---

## syslog, openlog, closelog, or setlogmask Subroutine

### Purpose

Controls the system log.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <syslog.h>
```

```
void openlog ( ID, LogOption, Facility)
const char *ID;
int LogOption, Facility;
```

```
void syslog ( Priority, Value,... )
int Priority;
const char *Value;
```

```
void closelog ( )
```

```
int setlogmask( MaskPriority)
int MaskPriority;
```

```
void bsdlog (Priority, Value,...)
int Priority;
const char *Value;
```

## Description

**Attention:** Do not use the **syslog**, **openlog**, **closelog**, or **setlogmask** subroutine in a multithreaded environment. See the multithread alternatives in the **syslog\_r** (“**syslog\_r**, **openlog\_r**, **closelog\_r**, or **setlogmask\_r** Subroutine” on page 332), **openlog\_r**, **closelog\_r**, or **setlogmask\_r** subroutine article. The **syslog** subroutine is not threadsafe; for threadsafe programs the **syslog\_r** subroutine should be used instead.

The **syslog** subroutine writes messages onto the system log maintained by the **syslogd** command.

**Note:** Messages passed to **syslog** that are longer than 900 bytes may be truncated by **syslogd** before being logged.

The message is similar to the **printf** *fmt* string, with the difference that *%m* is replaced by the current error message obtained from the **errno** global variable. A trailing new-line can be added to the message if needed.

Messages are read by the **syslogd** command and written to the system console or log file, or forwarded to the **syslogd** command on the appropriate host.

If special processing is required, the **openlog** subroutine can be used to initialize the log file.

Messages are tagged with codes indicating the type of *Priority* for each. A *Priority* is encoded as a *Facility*, which describes the part of the system generating the message, and as a level, which indicates the severity of the message.

If the **syslog** subroutine cannot pass the message to the **syslogd** command, it writes the message on the **/dev/console** file, provided the **LOG\_CONS** option is set.

The **closelog** subroutine closes the log file.

The **setlogmask** subroutine uses the bit mask in the *MaskPriority* parameter to set the new log priority mask and returns the previous mask.

The **LOG\_MASK** and **LOG\_UPTO** macros in the **sys/syslog.h** file are used to create the priority mask. Calls to the **syslog** subroutine with a priority mask that does not allow logging of that particular level of message causes the subroutine to return without logging the message.

## Parameters

*ID*                      Contains a string that is attached to the beginning of every message. The *Facility* parameter encodes a default facility from the previous list to be assigned to messages that do not have an explicit facility encoded.

## *LogOption*

Specifies a bit field that indicates logging options. The values of *LogOption* are:

### **LOG\_CONS**

Sends messages to the console if unable to send them to the **syslogd** command. This option is useful in daemon processes that have no controlling terminal.

### **LOG\_NDELAY**

Opens the connection to the **syslogd** command immediately, instead of when the first message is logged. This option is useful for programs that need to manage the order in which file descriptors are allocated.

### **LOG\_NOWAIT**

Logs messages to the console without waiting for forked children. Use this option for processes that enable notification of child termination through **SIGCHLD**; otherwise, the **syslog** subroutine may block, waiting for a child process whose exit status has already been collected.

### **LOG\_ODELAY**

Delays opening until the **syslog** subroutine is called.

### **LOG\_PID**

Logs the process ID with each message. This option is useful for identifying daemons.

## *Facility*

Specifies which of the following values generated the message:

### **LOG\_AUTH**

Indicates the security authorization system: the **login** command, the **su** command, and so on.

### **LOG\_DAEMON**

Logs system daemons.

### **LOG\_KERN**

Logs messages generated by the kernel. Kernel processes should use the **bsdlog** routine to generate **syslog** messages. The syntax of **bsdlog** is identical to **syslog**. The **bsdlog** messages can only be created by kernel processes and must be of **LOG\_KERN** priority. The **syslog** subroutine cannot log **LOG\_KERN** facility messages. Instead it will log **LOG\_USER** facility messages.

### **LOG\_LPR**

Logs the line printer spooling system.

### **LOG\_LOCAL0 through LOG\_LOCAL7**

Reserved for local use.

### **LOG\_MAIL**

Logs the mail system.

### **LOG\_NEWS**

Logs the news subsystem.

### **LOG\_UUCP**

Logs the UUCP subsystem.

### **LOG\_USER**

Logs messages generated by user processes. This is the default facility when none is specified.

<i>Priority</i>	Specifies the part of the system generating the message, and as a level, indicates the severity of the message. The level of severity is selected from the following list: <b>LOG_ALERT</b> Indicates a condition that should be corrected immediately; for example, a corrupted database. <b>LOG_CRIT</b> Indicates critical conditions; for example, hard device errors. <b>LOG_DEBUG</b> Displays messages containing information useful to debug a program. <b>LOG_EMERG</b> Indicates a panic condition reported to all users; system is unusable. <b>LOG_ERR</b> Indicated error conditions. <b>LOG_INFO</b> Indicates general information messages. <b>LOG_NOTICE</b> Indicates a condition requiring special handling, but not an error condition. <b>LOG_WARNING</b> Logs warning messages.
<i>MaskPriority</i>	Enables logging for the levels indicated by the bits in the mask that are set and disabled where the bits are not set. The default mask allows all priorities to be logged.
<i>Value</i>	Specifies the values given in the <i>Value</i> parameters and follows the the same syntax as the <b>printf</b> subroutine <i>Format</i> parameter.

## Examples

1. To log an error message concerning a possible security breach, such as the following, enter:  

```
syslog (LOG_ALERT, "who:internal error 23");
```
2. To initialize the log file, set the log priority mask, and log an error message, enter:  

```
openlog ("ftpd", LOG_PID, LOG_DAEMON);
setlogmask (LOG_UPTO (LOG_ERR));
syslog (LOG_INFO);
```
3. To log an error message from the system, enter:  

```
syslog (LOG_INFO | LOG_LOCAL2, "foobar error: %m");
```

## Related Information

The **profil** subroutine.

The **prof** command.

The **syslogd** daemon.

**\_end**, **\_etext**, or **edata** identifiers.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## syslog\_r, openlog\_r, closelog\_r, or setlogmask\_r Subroutine

### Purpose

Controls the system log.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <syslog.h>

int syslog_r (Priority, SysLogData, Format, . . .)
int Priority;
struct syslog_data * SysLogData;
const char * Format;

int openlog_r (ID, LogOption, Facility, SysLogData)
const char * ID;
int LogOption;
int Facility;

struct syslog_data *SysLogData;
void closelog_r (SysLogData)
struct syslog_data *SysLogData;

int setlogmask_r ( MaskPriority, SysLogData)
int MaskPriority;
struct syslog_data *SysLogData;
```

## Description

The **syslog\_r** subroutine writes messages onto the system log maintained by the **syslogd** daemon.

The messages are similar to the *Format* parameter in the **printf** subroutine, except that the %m field is replaced by the current error message obtained from the **errno** global variable. A trailing new-line character can be added to the message if needed.

Messages are read by the **syslogd** daemon and written to the system console or log file, or forwarded to the **syslogd** daemon on the appropriate host.

If a program requires special processing, you can use the **openlog\_r** subroutine to initialize the log file.

The **syslog\_r** subroutine takes as a second parameter a variable of the type **struct syslog\_data**, which should be provided by the caller. When that variable is declared, it should be set to the **SYSLOG\_DATA\_INIT** value, which specifies an initialization macro defined in the **sys/syslog.h** file. Without initialization, the data structure used to support the thread safety is not set up and the **syslog\_r** subroutine does not work properly.

Messages are tagged with codes indicating the type of *Priority* for each. A *Priority* is encoded as a *Facility*, which describes the part of the system generating the message, and as a level, which indicates the severity of the message.

If the **syslog\_r** subroutine cannot pass the message to the **syslogd** daemon, it writes the message the **/dev/console** file, provided the **LOG\_CONS** option is set.

The **closelog\_r** subroutine closes the log file.

The **setlogmask\_r** subroutine uses the bit mask in the *MaskPriority* parameter to set the new log priority mask and returns the previous mask.

The **LOG\_MASK** and **LOG\_UPTO** macros in the **sys/syslog.h** file are used to create the priority mask. Calls to the **syslog\_r** subroutine with a priority mask that does not allow logging of that particular level of message causes the subroutine to return without logging the message.

Programs using this subroutine must link to the **libpthreads.a** library.

## Parameters

<i>Priority</i>	<p>Specifies the part of the system generating the message and indicates the level of severity of the message. The level of severity is selected from the following list:</p> <ul style="list-style-type: none"><li>• A condition that should be corrected immediately, such as a corrupted database.</li><li>• A critical condition, such as hard device errors.</li><li>• A message containing information useful to debug a program.</li><li>• A panic condition reported to all users, such as an unusable system.</li><li>• An error condition.</li><li>• A general information message.</li><li>• A condition requiring special handling, other than an error condition.</li><li>• A warning message.</li></ul>
<i>SysLogData</i>	<p>Specifies a structure that contains the following information:</p> <ul style="list-style-type: none"><li>• The file descriptor for the log file.</li><li>• The status bits for the log file.</li><li>• A string for tagging the log entry.</li><li>• The mask of priorities to be logged.</li><li>• The default facility code.</li><li>• The address of the local logger.</li></ul>
<i>Format</i>	<p>Specifies the format, given in the same format as for the <b>printf</b> subroutine.</p>
<i>ID</i>	<p>Contains a string attached to the beginning of every message. The Facility parameter encodes a default facility from the previous list to be assigned to messages that do not have an explicit facility encoded.</p>
<i>LogOption</i>	<p>Specifies a bit field that indicates logging options. The values of <i>LogOption</i> are:</p> <p><b>LOG_CONS</b> Sends messages to the console if unable to send them to the <b>syslogd</b> command. This option is useful in daemon processes that have no controlling terminal.</p> <p><b>LOG_NDELAY</b> Opens the connection to the <b>syslogd</b> command immediately, instead of when the first message is logged. This option is useful for programs that need to manage the order in which file descriptors are allocated.</p> <p><b>LOG_NOWAIT</b> Logs messages to the console without waiting for forked children. Use this option for processes that enable notification of child termination through <b>SIGCHLD</b>; otherwise, the <b>syslog</b> subroutine may block, waiting for a child process whose exit status has already been collected.</p> <p><b>LOG_ODELAY</b> Delays opening until the <b>syslog</b> subroutine is called.</p> <p><b>LOG_PID</b> Logs the process ID with each message. This option is useful for identifying daemons.</p>

## Facility

Specifies which of the following values generated the message:

### LOG\_AUTH

Indicates the security authorization system: the **login** command, the **su** command, and so on.

### LOG\_DAEMON

Logs system daemons.

### LOG\_KERN

Logs messages generated by the kernel. Kernel processes should use the **bsdlog** routine to generate **syslog** messages. The syntax of **bsdlog** is identical to **syslog**. The **bsdlog** messages can only be created by kernel processes and must be of **LOG\_KERN** priority.

### LOG\_LPR

Logs the line printer spooling system.

### LOG\_LOCAL0 through LOG\_LOCAL7

Reserved for local use.

### LOG\_MAIL

Logs the mail system.

### LOG\_NEWS

Logs the news subsystem.

### LOG\_UUCP

Logs the UUCP subsystem.

### LOG\_USER

Logs messages generated by user processes. This is the default facility when none is specified.

- Remote file systems, such as the Andrew File System (AFS).
- The UUCP subsystem.
- Messages generated by user processes. This is the default facility when none is specified.

## MaskPriority

Enables logging for the levels indicated by the bits in the mask that are set, and disables logging where the bits are not set. The default mask allows all priorities to be logged.

## Return Values

- 0** Indicates that the subroutine was successful.
- 1** Indicates that the subroutine was not successful.

## Examples

1. To log an error message concerning a possible security breach, enter:  

```
syslog_r (LOG_ALERT, syslog_data_struct, "%s", "who:internal error 23");
```
2. To initialize the log file, set the log priority mask, and log an error message, enter:  

```
openlog_r ("ftpd", LOG_PID, LOG_DAEMON, syslog_data_struct);  
setlogmask_r (LOG_UPTO (LOG_ERR), syslog_data_struct);  
syslog_r (LOG_INFO, syslog_data_struct, "");
```
3. To log an error message from the system, enter:  

```
syslog_r (LOG_INFO | LOG_LOCAL2, syslog_data_struct, "system error: %m");
```

## Related Information

The **prof** command.

The **syslogd** daemon.

The **printf**, **fprintf**, **sprintf**, **wsprintf**, **vprintf**, **vfprintf**, **vsprintf**, or **vwsprintf** subroutine.

Subroutines Overview and List of Multithread Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## sys\_parm Subroutine

### Purpose

Provides a service for examining or setting kernel run-time tunable parameters.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int sys_parm ( cmd, parmflag, parm)
int cmd;
int parmflag;
struct vario *parm;
```

### Description

The **sys\_parm** subroutine is used to query and/or customize run-time operating system parameters.

**Note:** This is a replacement service for **sysconfig** with respect to querying or changing information in the **var** structure.

The **sys\_parm** subroutine:

- Works on both 32 bit and 64 bit platforms
- Requires appropriate privilege for its use.

The following operations are supported:

**SYSP\_GET**

Returns a structure containing the current value of the specified run-time parameter found in the **var** structure.

**SYSP\_SET**

Sets the value of the specified run-time parameter.

The run-time parameters that can be returned or set are found in the **var** structure as defined in **var.h**

### Parameters

*cmd*

Specifies the **SYSP\_GET** or **SYSP\_SET** function.

*parmflag*

Specifies the parameter upon which the function will act.

*parm*

Points to the user specified structure from which or to which the system parameter value is copied. *parm* points to a structure of type **vario** as defined in **var.h**.

The **vario** structure is an abstraction of the various fields in the **var** structure for which each field is size invariant. The size of the data does not depend on the execution environment of the kernel being 32 or 64 bit or the calling application being 32 or 64 bit.

## Examples

1. To examine the value of `v.v_iostrun` (collect disk usage statistics).

```
#include <sys/var.h>
#include <stdio.h>
struct vario myvar;
rc=sys_parm(SYSP_GET,SYSP_V_IOSTRUN,);
if(rc==0)
    printf("v.v_iostrun is set to %d\n",myvar.v.v_iostrun.value);
```

2. To change the value of `v.v_iostrun` (collect disk usage statistics).

```
#include <sys/var.h>
#include <stdio.h>
struct vario myvar;
myvar.v.v_iostrun.value=0; /* initialize to false */
rc=sys_parm(SYSP_SET,SYSP_V_IOSTRUN,);
if(rc==0)
    printf("disk usage statistics are not being collected\n");
```

Other parameters may be examined or set by changing the `parmflag` parameter.

## Return Values

These operations return a value of 0 upon successful completion of the subroutine. Otherwise or a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

**EACCES**  
**EINVAL**

The calling process does not have the required privilege. One of the following is true:

- The command is neither **SYSP\_GET** nor **SYSP\_SET**
- *parmflag* is out of range of parameters defined in `var.h`
- The value specified in the *parm* parameter is not a valid value for the field indicated by the *parmflag* parameter.

**EFAULT**

An invalid address was specified by the *parm* parameter.

## File

`sys/var.h`                      Contains structure definitions.

## Related Information

The **SYS\_GETPARMS** (“SYS\_GETPARMS sysconfig Operation” on page 320) **sysconfig** Operation, and **SYS\_SETPARMS** (“SYS\_SETPARMS sysconfig Operation” on page 327) **sysconfig** Operation

---

## system Subroutine

### Purpose

Runs a shell command.

## Library

Standard C Library (**libc.a**)

## Syntax

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

```
int system ( String )  
const char *String;
```

## Description

The **system** subroutine passes the *String* parameter to the **sh** command as input. Then the **sh** command interprets the *String* parameter as a command and runs it.

The **system** subroutine calls the **fork** subroutine to create a child process that in turn uses the **exec I** subroutine to run the **/usr/bin/sh** command, which interprets the shell command contained in the *String* parameter. When invoked on the Trusted Path, the **system** subroutine runs the Trusted Path shell (**/usr/bin/tsh**). The current process waits until the shell has completed, then returns the exit status of the shell. The exit status of the shell is returned in the same manner as a call to the **wait** or **waitpid** subroutine, using the structures in the **sys/wait.h** file.

The **system** subroutine ignores the **SIGINT** and **SIGQUIT** signals, and blocks the **SIGCHILD** signal while waiting for the command specified by the *String* parameter to terminate. If this might cause the application to miss a signal that would have killed it, the application should use the value returned by the **system** subroutine to take the appropriate action if the command terminated due to receipt of a signal. The **system** subroutine does not affect the termination status of any child of the calling process unless that process was created by the **system** subroutine. The **system** subroutine does not return until the child process has terminated.

## Parameters

*String*        Specifies a valid **sh** shell command.

**Note:** The **system** subroutine runs only **sh** shell commands. The results are unpredictable if the *String* parameter is not a valid **sh** shell command.

## Return Values

Upon successful completion, the **system** subroutine returns the exit status of the shell. The exit status of the shell is returned in the same manner as a call to the **wait** or **waitpid** subroutine, using the structures in the **sys/wait.h** file.

If the *String* parameter is a null pointer and a command processor is available, the **system** subroutine returns a nonzero value. If the **fork** subroutine fails or if the exit status of the shell cannot be obtained, the **system** subroutine returns a value of -1. If the **exec I** subroutine fails, the **system** subroutine returns a value of 127. In all cases, the **errno** global variable is set to indicate the error.

## Error Codes

The **system** subroutine fails if any of the following are true:

- EAGAIN**     The system-imposed limit on the total number of running processes, either systemwide or by a single user ID, was exceeded.
- EINTR**      The **system** subroutine was interrupted by a signal that was caught before the requested process was started. The **EINTR** error code will never be returned after the requested process has begun.
- ENOMEM**    Insufficient storage space is available.

## Related Information

The **execl** subroutine, **exit** subroutine, **fork** subroutine, **pipe** subroutine, **wait** (“wait, waitpid, wait3, or wait364 Subroutine” on page 440) subroutine, **waitpid** (“wait, waitpid, wait3, or wait364 Subroutine” on page 440) subroutine.

The **sh** command.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tan, tanf, or tanl Subroutine

### Purpose

Computes the tangent.

### Syntax

```
#include <math.h>
```

```
float tanf (x)  
float x;
```

```
long double tanl (x)  
long double x;
```

```
double tan (x)  
double x;
```

### Description

The **tan**, **tanf**, and **tanl** subroutines compute the tangent of the *x* parameter, measured in radians.

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these functions. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

*x* Specifies the value to be computed.

### Return Values

Upon successful completion, the **tan**, **tanf**, and **tanl** subroutines return the tangent of *x*.

If the correct value would cause underflow, and is not representable, a range error may occur, and 0.0 is returned.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$ , *x* is returned.

If *x* is subnormal, a range error may occur and *x* should be returned.

If  $x$  is  $\pm\text{Inf}$ , a domain error occurs, and a NaN returned.

If the correct value would cause underflow, and is representable, a range error may occur and the correct value is returned.

If the correct value would cause overflow, a range error occurs and the **tan**, **tanf**, and **tanh** subroutines return the value of the macro **HUGE\_VAL**, **HUGE\_VALF**, and **HUGE\_VALL**, respectively.

## Error Codes

The **tan**, **tanf**, and **tanh** subroutines lose accuracy when passed a large value for the  $x$  parameter. Since the machine value of  $\pi$  can only approximate its infinitely precise value, the remainder of  $x/(2 * \pi)$  becomes less accurate as  $x$  becomes larger. Similar loss of accuracy occurs for the **tan**, **tanf**, and **tanh** subroutines during argument reduction of large arguments.

## Related Information

atanf or atanl Subroutine, feclearexcept Subroutine, fetestexcept Subroutine, and class, \_class, finite, isnan, or unordered Subroutines in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

**math.h** in *AIX 5L Version 5.2 Files Reference*.

---

## tanh, tanhf, or tanhl Subroutine

### Purpose

Computes the hyperbolic tangent.

### Syntax

```
#include <math.h>
```

```
float tanhf (x)  
float x;
```

```
long double tanhl (x)  
double x;
```

```
double tanh (x)  
double x;
```

### Description

The **tanhf**, **tanh**, and **tanhl** subroutines compute the hyperbolic tangent of the  $x$ .

An application wishing to check for error situations should set the **errno** global variable to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

### Parameters

$x$  Specifies the value to be computed.

## Return Values

Upon successful completion, the **tanhf**, **tanhl**, and **tanh** subroutines return the hyperbolic tangent of *x*.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$ , *x* is returned.

If *x* is  $\pm \text{Inf}$ ,  $\pm 1$  is returned.

If *x* is subnormal, a range error may occur and *x* should be returned.

## Related Information

The “sin, sinh, or sinl Subroutine” on page 193.

atanf or atanl Subroutine, feclearexcept Subroutine, fetestexcept Subroutine, and class, \_class, finite, isnan, or unordered Subroutines in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

**math.h** in *AIX 5L Version 5.2 Files Reference*.

---

## tcb Subroutine

### Purpose

Alters the Trusted Computing Base (TCB) status of a file.

### Library

Security Library (**libc.a**)

### Syntax

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

```
int tcb ( Path, Flag)
```

```
char *Path;
```

```
int Flag;
```

### Description

The **tcb** subroutine provides a mechanism to query or set the TCB attributes of a file.

This subroutine is not safe for use with multiple threads. To call this subroutine from a threaded application, enclose the call with the **\_libs\_rmutex** lock. See “Making a Subroutine Safe for Multiple Threads” in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs* for more information about this lock.

### Parameters

*Path* Specifies the path name of the file whose TCB status is to be changed.

*Flag* Specifies the function to be performed. Valid values are defined in the **sys/tcb.h** file and include the following:

**TCB\_ON**

Enables the TCB attribute of a file.

**TCB\_OFF**

Disables the Trusted Process and TCB attributes of a file.

**TCB\_QUERY**

Queries the TCB status of a file. This function returns one of the preceding values.

## Return Values

Upon successful completion, the **tcb** subroutine returns a value of 0 if the *Flags* parameter is either **TCB\_ON** or **TCB\_OFF**. If the *Flags* parameter is **TCB\_QUERY**, the current status is returned. If the **tcb** subroutine fails, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **tcb** subroutine fails if one of the following is true:

**EINVAL** The *Flags* parameter is not one of **TCB\_ON**, **TCB\_OFF**, or **TCB\_QUERY**.

**EPERM** Not authorized to perform this operation.

**ENOENT** The file specified by the *Path* parameter does not exist.

**EROFS** The file system is read-only.

**EBUSY** The file specified by the *Path* parameter is currently open for writing.

**EACCES** Access permission is denied for the file specified by the *Path* parameter.

## Security

Access Control: The calling process must have search permission for the object named by the *Path* parameter. Only the root user can set the **tcb** attributes of a file.

## Related Information

The **chmod** or **fchmod** subroutine, **statx**, **stat**, **lstat**, **fstatx**, **fstat**, **fullstat**, or **ffullstat** (“statx, stat, lstat, fstatx, fstat, fullstat, ffullstat, stat64, lstat64, or fstat64 Subroutine” on page 277) subroutine.

The **chmod** command.

List of Security and Auditing Subroutines, Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tcdrain Subroutine

### Purpose

Waits for output to complete.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <termios.h>
```

```
int tcdrain( FileDescriptor)
int FileDescriptor;
```

## Description

The **tcdrain** subroutine waits until all output written to the object referred to by the *FileDescriptor* parameter has been transmitted.

## Parameter

*FileDescriptor* Specifies an open file descriptor.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **tcdrain** subroutine is unsuccessful if one of the following is true:

**EBADF** The *FileDescriptor* parameter does not specify a valid file descriptor.  
**EINTR** A signal interrupted the **tcdrain** subroutine.  
**EIO** The process group of the writing process is orphaned, and the writing process does not ignore or block the **SIGTTOU** signal.  
**ENOTTY** The file associated with the *FileDescriptor* parameter is not a terminal.

## Example

To wait until all output has been transmitted, enter:

```
rc = tcdrain(stdout);
```

## Related Information

The **tcflow** (“tcflow Subroutine”) subroutine, **tcflush** (“tcflush Subroutine” on page 344) subroutine, **tcsendbreak** (“tcsendbreak Subroutine” on page 348) subroutine.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tcflow Subroutine

### Purpose

Performs flow control functions.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <termios.h>
```

```
int tcflow( FileDescriptor, Action)
int FileDescriptor;
int Action;
```

## Description

The **tcflow** subroutine suspends transmission or reception of data on the object referred to by the *FileDescriptor* parameter, depending on the value of the *Action* parameter.

## Parameters

<i>FileDescriptor</i>	Specifies an open file descriptor.
<i>Action</i>	Specifies one of the following:
<b>TCOOFF</b>	Suspend output.
<b>TCOON</b>	Restart suspended output.
<b>TCIOFF</b>	Transmit a STOP character, which is intended to cause the terminal device to stop transmitting data to the system. See the description of IXOFF in the Input Modes section of the <b>termios.h</b> file.
<b>TCION</b>	Transmit a START character, which is intended to cause the terminal device to start transmitting data to the system. See the description of IXOFF in the Input Modes section of the <b>termios.h</b> file.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **tcflow** subroutine is unsuccessful if one of the following is true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter does not specify a valid file descriptor.
<b>EINVAL</b>	The <i>Action</i> parameter does not specify a proper value.
<b>EIO</b>	The process group of the writing process is orphaned, and the writing process does not ignore or block the <b>SIGTTOU</b> signal.
<b>ENOTTY</b>	The file associated with the <i>FileDescriptor</i> parameter is not a terminal.

## Example

To restart output from a terminal device, enter:

```
rc = tcflow(stdout, TCION);
```

## Related Information

The **tcdrain** (“tcdrain Subroutine” on page 342) subroutine, **tcflush** (“tcflush Subroutine”) subroutine, **tcsendbreak** (“tcsendbreak Subroutine” on page 348) subroutine.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tcflush Subroutine

### Purpose

Discards data from the specified queue.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <termios.h>
```

```
int tcflush( FileDescriptor, QueueSelector )  
int FileDescriptor;  
int QueueSelector;
```

## Description

The **tcflush** subroutine discards any data written to the object referred to by the *FileDescriptor* parameter, or data received but not read by the object referred to by *FileDescriptor*, depending on the value of the *QueueSelector* parameter.

## Parameters

*FileDescriptor*

Specifies an open file descriptor.

*QueueSelector*

Specifies one of the following:

### TCIFLUSH

Flush data received but not read.

### TCOFLUSH

Flush data written but not transmitted.

### TCIOFLUSH

Flush both of the following:

- Data received but not read
- Data written but not transmitted

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **tcflush** subroutine is unsuccessful if one of the following is true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter does not specify a valid file descriptor.
<b>EINVAL</b>	The <i>QueueSelector</i> parameter does not specify a proper value.
<b>EIO</b>	The process group of the writing process is orphaned, and the writing process does not ignore or block the <b>SIGTTOU</b> signal.
<b>ENOTTY</b>	The file associated with the <i>FileDescriptor</i> parameter is not a terminal.

## Example

To flush the output queue, enter:

```
rc = tcflush(2, TCOFLUSH);
```

## Related Information

The **tcdrain** (“tcdrain Subroutine” on page 342) subroutine, **tcflow** (“tcflow Subroutine” on page 343) subroutine, **tcsendbreak** (“tcsendbreak Subroutine” on page 348) subroutine.

## tcgetattr Subroutine

### Purpose

Gets terminal state.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <termios.h>
```

```
int tcgetattr ( FileDescriptor, TermiosPointer )
int FileDescriptor;
struct termios *TermiosPointer;
```

### Description

The **tcgetattr** subroutine gets the parameters associated with the object referred to by the *FileDescriptor* parameter and stores them in the **termios** structure referenced by the *TermiosPointer* parameter. This subroutine is allowed from a background process; however, the terminal attributes may subsequently be changed by a foreground process.

Whether or not the terminal device supports differing input and output baud rates, the baud rates stored in the **termios** structure returned by the **tcgetattr** subroutine reflect the actual baud rates, even if they are equal.

**Note:** If differing baud rates are not supported, returning a value of 0 as the input baud rate is obsolete.

### Parameters

<i>FileDescriptor</i>	Specifies an open file descriptor.
<i>TermiosPointer</i>	Points to a <b>termios</b> structure.

### Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **tcgetattr** subroutine is unsuccessful if one of the following is true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter does not specify a valid file descriptor.
<b>ENOTTY</b>	The file associated with the <i>FileDescriptor</i> parameter is not a terminal.

### Examples

To get the current terminal state information, enter:

```
rc = tcgetattr(stdout, &my_termios);
```

## Related Information

The **tcsetattr** (“tcsetattr Subroutine” on page 349) subroutine.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tcgetpgrp Subroutine

### Purpose

Gets foreground process group ID.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
pid_t tcgetpgrp ( FileDescriptor )  
int FileDescriptor;
```

### Description

The **tcgetpgrp** subroutine returns the value of the process group ID of the foreground process group associated with the terminal. The function can be called from a background process; however, the foreground process can subsequently change the information.

### Parameters

*FileDescriptor*                      Indicates the open file descriptor for the terminal special file.

### Return Values

Upon successful completion, the process group ID of the foreground process is returned. If there is no foreground process group, a value greater than 1 that does not match the process group ID of any existing process group is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **tcgetpgrp** subroutine is unsuccessful if one of the following is true:

- EBADF**      The *FileDescriptor* argument is not a valid file descriptor.
- EINVAL**     The function is not appropriate for the file associated with the *FileDescriptor* argument.
- ENOTTY**    The calling process does not have a controlling terminal or the file is not the controlling terminal.

## Related Information

The **setpgid** (“setpgid or setpgrp Subroutine” on page 146) subroutine, **setsid** (“setsid Subroutine” on page 150) subroutine, **tcsetpgrp** (“tcsetpgrp Subroutine” on page 351) subroutine.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tcsendbreak Subroutine

### Purpose

Sends a break on an asynchronous serial data line.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <termios.h>
```

```
int tcsendbreak( FileDescriptor, Duration )
int FileDescriptor;
int Duration;
```

### Description

If the terminal is using asynchronous serial data transmission, the **tcsendbreak** subroutine causes transmission of a continuous stream of zero-valued bits for a specific duration.

If the terminal is not using asynchronous serial data transmission, the **tcsendbreak** subroutine returns without taking any action.

Pseudo-terminals and LFT do not generate a break condition. They return without taking any action.

### Parameters

<i>FileDescriptor</i>	Specifies an open file descriptor.
<i>Duration</i>	Specifies the number of milliseconds that zero-valued bits are transmitted. If the value of the <i>Duration</i> parameter is 0, it causes transmission of zero-valued bits for at least 250 milliseconds and not longer than 500 milliseconds. If <i>Duration</i> is not 0, it sends zero-valued bits for <i>Duration</i> milliseconds.

### Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

The **tcsendbreak** subroutine is unsuccessful if one or both of the following are true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter does not specify a valid open file descriptor.
<b>EIO</b>	The process group of the writing process is orphaned, and the writing process does not ignore or block the <b>SIGTTOU</b> signal.
<b>ENOTTY</b>	The file associated with the <i>FileDescriptor</i> parameter is not a terminal.

### Examples

1. To send a break condition for 500 milliseconds, enter:  

```
rc = tcsendbreak(stdout,500);
```
2. To send a break condition for 25 milliseconds, enter:  

```
rc = tcsendbreak(1,25);
```

This could also be performed using the default *Duration* by entering:

```
rc = tcsendbreak(1, 0);
```

## Related Information

The **tcdrain** (“tcdrain Subroutine” on page 342) subroutine, **tcflow** (“tcflow Subroutine” on page 343) subroutine, **tcflush** (“tcflush Subroutine” on page 344) subroutine.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tcsetattr Subroutine

### Purpose

Sets terminal state.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <termios.h>
```

```
int tcsetattr (FileDescriptor, OptionalActions, TermiosPointer)
int FileDescriptor, OptionalActions;
const struct termios * TermiosPointer;
```

### Description

The **tcsetattr** subroutine sets the parameters associated with the object referred to by the *FileDescriptor* parameter (unless support required from the underlying hardware is unavailable), from the **termios** structure referenced by the *TermiosPointer* parameter.

The value of the *OptionalActions* parameter determines how the **tcsetattr** subroutine is handled.

The 0 baud rate (B0) is used to terminate the connection. If B0 is specified as the output baud rate when the **tcsetattr** subroutine is called, the modem control lines are no longer asserted. Normally, this disconnects the line.

Using 0 as the input baud rate in the **termios** structure to cause **tcsetattr** to change the input baud rate to the same value as that specified by the value of the output baud rate, is obsolete.

If an attempt is made using the **tcsetattr** subroutine to set:

- An unsupported baud rate
- Baud rates, such that the input and output baud rates differ and the hardware does not support that combination
- Other features not supported by the hardware

but the **tcsetattr** subroutine is able to perform some of the requested actions, then the subroutine returns successfully, having set all supported attributes and leaving the above unsupported attributes unchanged.

If no part of the request can be honored, the **tcsetattr** subroutine returns a value of -1 and the **errno** global variable is set to **EINVAL**.

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 the **tcgetattr** subroutine returns the actual state of the terminal device (reflecting both the changes made and not made in the previous **tcsetattr** call). The **tcsetattr** subroutine does not change the values in the **termios** structure whether or not it actually accepts them.

If the **tcsetattr** subroutine is called by a process which is a member of a background process group on a *FileDescriptor* associated with its controlling terminal, a **SIGTTOU** signal is sent to the background process group. If the calling process is blocking or ignoring **SIGTTOU** signals, the process performs the operation and no signal is sent.

## Parameters

<i>FileDescriptor</i>	Specifies an open file descriptor.
<i>OptionalActions</i>	Specifies one of the following values: <b>TCSANOW</b> The change occurs immediately. <b>TCSADRAIN</b> The change occurs after all output written to the object referred to by <i>FileDescriptor</i> has been transmitted. This function should be used when changing parameters that affect output. <b>TCSAFLUSH</b> The change occurs after all output written to the object referred to by <i>FileDescriptor</i> has been transmitted. All input that has been received but not read is discarded before the change is made.
<i>TermiosPointer</i>	Points to a <b>termios</b> structure.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **tcsetattr** subroutine is unsuccessful if one of the following is true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter does not specify a valid file descriptor.
<b>EINTR</b>	A signal interrupted the <b>tcsetattr</b> subroutine.
<b>EINVAL</b>	The <i>OptionalActions</i> argument is not a proper value, or an attempt was made to change an attribute represented in the <b>termios</b> structure to an unsupported value.
<b>EIO</b>	The process group of the writing process is orphaned, and the writing process does not ignore or block the <b>SIGTTOU</b> signal.
<b>ENOTTY</b>	The file associated with the <i>FileDescriptor</i> parameter is not a terminal.

## Example

To set the terminal state after the current output completes, enter:

```
rc = tcsetattr(stdout, TCSADRAIN, &my_termios);
```

## Related Information

The **cfgetispeed** subroutine, **tcgetattr** (“tcgetattr Subroutine” on page 346) subroutine.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tcsetpgrp Subroutine

### Purpose

Sets foreground process group ID.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int tcsetpgrp ( FileDescriptor, ProcessGroupID )
int FileDescriptor;
pid_t ProcessGroupID;
```

### Description

If the process has a controlling terminal, the **tcsetpgrp** subroutine sets the foreground process group ID associated with the terminal to the value of the *ProcessGroupID* parameter. The file associated with the *FileDescriptor* parameter 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 the *ProcessGroupID* parameter must match a process group ID of a process in the same session as the calling process.

### Parameters

<i>FileDescriptor</i>	Specifies an open file descriptor.
<i>ProcessGroupID</i>	Specifies the process group identifier.

### Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

### Error Codes

This function is unsuccessful if one of the following is true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter is not a valid file descriptor.
<b>EINVAL</b>	The <i>ProcessGroupID</i> parameter is invalid.
<b>ENOTTY</b>	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.
<b>EPERM</b>	The <i>ProcessGroupID</i> parameter is valid, but does not match the process group ID of a process in the same session as the calling process.

### Related Information

The **tcgetpgrp** (“tcgetpgrp Subroutine” on page 347) subroutine.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## termdef Subroutine

### Purpose

Queries terminal characteristics.

### Library

Standard C Library (**libc.a**)

### Syntax

```
char *termdef ( FileDescriptor, Characteristic )
int FileDescriptor;
char Characteristic;
```

### Description

The **termdef** subroutine returns a pointer to a null-terminated, static character string that contains the value of a characteristic defined for the terminal specified by the *FileDescriptor* parameter.

### Asynchronous Terminal Support

Shell profiles usually set the **TERM** environment variable each time you log in. The **stty** command allows you to change the lines and columns (by using the *lines* and *cols* options). This is preferred over changing the **LINES** and **COLUMNS** environment variables, since the **termdef** subroutine examines the environment variables last. You consider setting **LINES** and **COLUMNS** environment variables if:

- You are using an asynchronous terminal and want to override the *lines* and *cols* setting in the **terminfo** database  
OR
- Your asynchronous terminal has an unusual number of lines or columns and you are running an application that uses the **termdef** subroutine but not an application which uses the **terminfo** database (for example, **curses**).

This is because the **curses** initialization subroutine, **setupterm** (“setupterm Subroutine” on page 624), calls the **termdef** subroutine to determine the number of lines and columns on the display. If the **termdef** subroutine cannot supply this information, the **setupterm** subroutine uses the values in the **terminfo** database.

### Parameters

*FileDescriptor*                      Specifies an open file descriptor.

*Characteristic*

Specifies the characteristic that is to be queried. The following values can be specified:

- c** Causes the **termdef** subroutine to query for the number of "columns" for the terminal. This is determined by performing the following actions:
  1. It requests a copy of the terminal's **winsize** structure by issuing the **TIOCGWINSZ ioctl**. If **ws\_col** is not 0, the **ws\_col** value is used.
  2. If the **TIOCGWINSZ ioctl** is unsuccessful or if **ws\_col** is 0, the **termdef** subroutine attempts to use the value of the **COLUMNS** environment variable.
  3. If the **COLUMNS** environment variable is not set, the **termdef** subroutine returns a pointer to a null string.
  
- l** Causes the **termdef** subroutine to query for the number of "lines" (or rows) for the terminal. This is determined by performing the following actions:
  1. It requests a copy of the terminal's **winsize** structure by issuing the **TIOCGWINSZ ioctl**. If **ws\_row** is not 0, the **ws\_row** value is used.
  2. If the **TIOCGWINSZ ioctl** is unsuccessful or if **ws\_row** is 0, the **termdef** subroutine attempts to use the value of the **LINES** environment variable.
  3. If the **LINES** environment variable is not set, the **termdef** subroutine returns a pointer to a null string.

**Characters other than c or l**

Cause the **termdef** subroutine to query for the "terminal type" of the terminal. This is determined by performing the following actions:

1. The **termdef** subroutine attempts to use the value of the **TERM** environment variable.
2. If the **TERM** environment variable is not set, the **termdef** subroutine returns a pointer to string set to "dumb".

## Examples

1. To display the terminal type of the standard input device, enter:

```
printf("%s\n", termdef(0, 't'));
```

2. To display the current lines and columns of the standard output device, enter:

```
printf("lines\tcolumns\n%s\t%s\n", termdef(2, 'l'),  
      termdef(2, 'c'));
```

**Note:** If the **termdef** subroutine is unable to determine a value for lines or columns, it returns pointers to null strings.

## Related Information

The **setupterm** ("setupterm Subroutine" on page 624) subroutine.

The **stty** command.

The Input and Output Handling Programmer's Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## test\_and\_set Subroutine

### Purpose

Atomically tests and sets a memory location.

## Library

Standard C library (**libc.a**)

## Syntax

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

boolean_t test_and_set (word_addr, mask)
atomic_p word_addr;
int mask;
```

## Description

The **test\_and\_set** subroutine attempts to atomically OR the value stored at *word\_addr* with the value specified by *mask*. If any bit in *mask* was already set in the value stored at *word\_addr*, no update is made.

## Parameters

<i>word_addr</i>	Specifies the address of the memory location to be set.
<i>mask</i>	Specifies the mask value to be used to set the memory location specified by <i>word_addr</i> .

## Return Values

The **test\_and\_set** subroutine returns true if the the value stored at *word\_addr* was updated. Otherwise, it returns false.

## Related Information

The **fetch\_and\_and** or **fetch\_and\_or** Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

---

## tgamma, tgammaf, or tgamma1 Subroutine

### Purpose

Computes the gamma.

### Syntax

```
#include <math.h>

double tgamma (x)
double x;

float tgammaf (x)
float x;

long double tgamma1 (x)
long double x;
```

### Description

The **tgamma**, **tgammaf**, and **tgamma1** subroutines compute the **gamma** function of *x*.

An application wishing to check for error situations should set **errno** to zero and call **feclearexcept(FE\_ALL\_EXCEPT)** before calling these subroutines. Upon return, if **errno** is nonzero or **fetestexcept(FE\_INVALID | FE\_DIVBYZERO | FE\_OVERFLOW | FE\_UNDERFLOW)** is nonzero, an error has occurred.

## Parameters

*x* Specifies the value to be computed.

## Return Values

Upon successful completion, the **tgamma**, **tgammaf**, and **tgammal** subroutines return **Gamma(x)**.

If *x* is a negative integer, a domain error occurs, and either a NaN (if supported), or an implementation-defined value is returned.

If the correct value would cause overflow, a range error occurs and the **tgamma**, **tgammaf**, and **tgammal** subroutines return the value of the macro **HUGE\_VAL**, **HUGE\_VALF**, or **HUGE\_VALL**, respectively.

If *x* is NaN, a NaN is returned.

If *x* is +Inf, *x* is returned.

If *x* is ±0, a pole error occurs, and the **tgamma**, **tgammaf**, and **tgammal** subroutines return ±**HUGE\_VAL**, ±**HUGE\_VALF**, and ±**HUGE\_VALL**, respectively.

If *x* is -Inf, a domain error occurs, and either a NaN (if supported), or an implementation-defined value is returned.

## Related Information

feclearexcept Subroutine, fetestexcept Subroutine, and lgamma, lgammal, or gamma Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

**math.h** in *AIX 5L Version 5.2 Files Reference*.

---

## times Subroutine

### Purpose

Gets process and waited-for child process times

### Syntax

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

```
clock_t times (buffer)
struct tms *buffer;
```

### Description

The **times** subroutine fills the **tms** structure pointed to by *buffer* with time-accounting information. The **tms** structure is defined in **<sys/times.h>**.

All times are measured in terms of the number of clock ticks used.

The times of a terminated child process is included in the *tms\_cutime* and *tms\_cstime* elements of the parent when the **wait** or **waitpid** subroutine returns the process ID of the terminated child. If a child process has not waited for its children, their times are not included in its times.

- The **tms\_utime** structure member is the CPU time charged for the execution of user instructions of the calling process.

- The **tms\_stime** structure member is the CPU time charged for execution by the system on behalf of the calling process.
- The **tms\_cutime** structure member is the sum of the **tms\_utime** and **tms\_cutime** times of the child processes.
- The **tms\_cstime** structure member is the sum of the **tms\_stime** and **tms\_cstime** times of the child processes.

Applications should use `sysconf(_SC_CLK_TCK)` to determine the number of clock ticks per second as it may vary from system to system.

## Parameters

*\*buffer* Points to the **tms** structure.

## Return Values

Upon successful completion, the **times** subroutine returns the elapsed real time, in clock ticks, since an arbitrary point in the past (for example, system startup time). This point does not change from one invocation of the **times** subroutine within the process to another. The return value may overflow the possible range of type *clock\_t*. If the **times** subroutine fails, `(clock_t)-1` is returned, and the **errno** global variable is set to indicate the error.

## Examples

### Timing a Database Lookup

The following example defines two functions, **start\_clock** and **end\_clock**, that are used to time a lookup. It also defines variables of type **clock\_t** and **tms** to measure the duration of transactions. The **start\_clock** function saves the beginning times given by the **times** subroutine. The **end\_clock** function gets the ending times and prints the difference between the two times.

```
#include <sys/times.h>
#include <stdio.h>
...
void start_clock(void);
void end_clock(char *msg);
...
static clock_t st_time;
static clock_t en_time;
static struct tms st_cpu;
static struct tms en_cpu;
...
void
start_clock()
{
    st_time = times(&st_cpu);
}

/* This example assumes that the result of each subtraction is within the range of values that can
   be represented in an integer type. */
void
end_clock(char *msg)
{
    en_time = times(&en_cpu);

    fputs(msg,stdout);
    printf("Real Time: %jd, User Time %jd, System Time %jd\n",
        (intmax_t)(en_time - st_time),
        (intmax_t)(en_cpu.tms_utime - st_cpu.tms_utime),
        (intmax_t)(en_cpu.tms_stime - st_cpu.tms_stime));
}
}
```

## Related Information

“sysconf Subroutine” on page 311 and “wait, waitpid, wait3, or wait364 Subroutine” on page 440

The gettimer, settimer, restimer, stime, or time Subroutine, getinterval, incinterval, absinterval, resinc, resabs, alarm, ualarm, getitimer or setitimer Subroutine, exec: execl, execl, execlp, execv, execve, execcvp, or exect Subroutine, and fork, f\_fork, or vfork Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

---

## timezone Subroutine

**Attention:** Do not use the **tzset** subroutine, from **libc.a**, when linking **libc.a libbsd.a**. The **tzset** subroutine uses the global external variable **timezone** which conflicts with the **timezone** subroutine in **libbsd.a**. This name collision can cause unpredictable results.

## Purpose

Returns the name of the timezone associated with the first argument.

## Library

Berkeley compatibility library (**libbsd.a**) (for **timezone** only)

## Syntax

```
#include <time.h>
char *timezone(zone, dst)
int zone;
int dst;

#include <time.h>
#include <limits.h>
int zone;
int dst;
char czone[TZNAME_MAX+1];
```

## Description

The **timezone** subroutine returns the name of the timezone associated with the first argument which is measured in minutes westward from Greenwich. If the environment variable **TZ** is set, the first argument is ignored and the current timezone is calculated from the value of **TZ**. If the second argument is 0, the standard name is returned otherwise the Daylight Saving Time name is returned. If **TZ** is not set, then the internal table is searched for a matching timezone. If the timezone does not appear in the built in table then difference from GMT is produced.

Timezone returns a pointer to static data that will be overwritten by subsequent calls.

## Parameters

<i>zone</i>	Specifies minutes westward from Greenwich.
<i>dst</i>	Specifies whether to return Standard time or Daylight Savings time.
<i>czone</i>	Specifies a buffer of size <b>TZNAME_MAX+1</b> , that the result is placed in.

## Return Values

**timezone** returns a pointer to static data that contains the name of the timezone.

## Errors

There are no errors defined.

## Related Information

Subroutines Overview

List of Multi-threaded Programming Subroutines

---

## thread\_post Subroutine

### Purpose

Posts a thread of an event completion.

### Library

Standard C library (**libc.a**)

### Syntax

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

```
int thread_post( tid)
tid_t tid;
```

### Description

The **thread\_post** subroutine posts the thread whose thread ID is indicated by the value of the *tid* parameter, of the occurrence of an event. If the posted thread is waiting in **thread\_wait**, it will be awakened immediately. If it not waiting in **thread\_wait**, the next call to **thread\_wait** does not block but returns with success immediately.

Multiple posts to the same thread without an intervening wait by the specified thread will only count as a single post. The posting remains in effect until the indicated thread calls the **thread\_wait** subroutine upon which the posting gets cleared.

The **thread\_wait** and the **thread\_post** subroutine can be used by applications to implement a fast IPC mechanism between threads in different processes.

### Parameters

*tid* Specifies the thread ID of the thread to be posted.

### Return Values

On successful completion, the **thread\_post** subroutine returns a value of **0**. If unsuccessful, a value of **-1** is returned and the global variable **errno** is set to indicate the error.

### Error Codes

<b>ESRCH</b>	This indicated thread is non-existent or the thread has exited or is exiting.
<b>EPERM</b>	The real or effective user ID does not match the real or effective user ID of the thread being posted, or else the calling process does not have root user authority.

## Related Information

The `thread_wait` (“thread\_wait Subroutine” on page 362) subroutine, and `thread_post_many` (“thread\_post\_many Subroutine”) subroutine.

---

## thread\_post\_many Subroutine

### Purpose

Posts one or more threads of an event completion.

### Library

Standard C library (**libc.a**)

### Syntax

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

```
int thread_post_many( nthreads, tidp, erridp)
int nthreads;
tid_t * tidp;
tid_t * erridp;
```

### Description

The `thread_post_many` subroutine posts one or more threads of the occurrence of the event. The number of threads to be posted is specified by the value of the `nthreads` parameter, while the `tidp` parameter points to an array of thread IDs of threads that need to be posted. The subroutine works just like the `thread_post` subroutine but can be used to post to multiple threads at the same time.

A maximum of 512 threads can be posted in one call to the `thread_post_many` subroutine.

An optional address to a thread ID field may be passed in the `erridp` parameter. This field is normally ignored by the kernel unless the subroutine fails because the calling process has no permissions to post to any one of the specified threads. In this case, the kernel posts all threads in the array pointed at by the `tidp` parameter up to the first failing thread and fills the `erridp` parameter with the failing thread's ID.

### Parameters

*nthreads*

Specifies the number of threads to be posted.

*tidp*

Specifies the address of an array of thread IDs corresponding to the list of threads to be posted.

*erridp*

Either NULL or specifies the pointer to a thread ID variable in which the kernel will return the thread ID of the first failing thread when an **errno** of **EPERM** is set.

### Return Values

On successful completion, the `thread_post_many` subroutine returns a value of **0**. If unsuccessful, a value of **-1** is returned and the global variable **errno** is set to indicate the error.

### Error Codes

The `thread_post_many` subroutine is unsuccessful when one of the following is true:

**ESRCH**

None of the indicated threads are existent or they have all exited or are exiting.

**EPERM**

The real or effective user ID does not match the real or effective user ID of one or more threads being posted, or else the calling process does not have root user authority.

**EFAULT**

The *tidp* parameter points to a location outside of the address space of the process.

**EINVAL**

A negative value or a value greater than 512 was specified in the *nthreads* parameter.

## Related Information

The **thread\_wait** (“thread\_wait Subroutine” on page 362) subroutine, and **thread\_post** (“thread\_post Subroutine” on page 358) subroutine.

---

## thread\_self Subroutine

### Purpose

Returns the caller’s kernel thread ID.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/thread.h>
tid_t thread_self ()
```

### Description

The **thread\_self** subroutine returns the caller’s kernel thread ID. The kernel thread ID may be useful for the **bindprocessor** and **ptrace** subroutines. The **ps**, **trace**, and **vmstat** commands also report kernel thread IDs, thus this subroutine can be useful for debugging multi-threaded programs.

The kernel thread ID is unrelated with the thread ID used in the threads library (**libpthreads.a**) and returned by the **pthread\_self** subroutine.

### Return Values

The **thread\_self** subroutine returns the caller’s kernel thread ID.

## Related Information

The **bindprocessor** subroutine, **pthread\_self** subroutine, **ptrace** subroutine.

---

## thread\_setsched Subroutine

### Purpose

Changes the scheduling policy and priority of a kernel thread.

### Library

Standard C library (**libc.a**)

### Syntax

```
#include <sys/sched.h>
#include <sys/pri.h>
```

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

```
int thread_setsched ( tid, priority, policy)
tid_t tid;
int priority;
int policy;
```

## Description

The **thread\_setsched** subroutine changes the scheduling policy and priority of a kernel thread. User threads (pthreads) have their own scheduling attributes that in some cases allow a pthread to execute on top of multiple kernel threads. Therefore, if the policy or priority change is being granted on behalf of a pthread, then the pthreads contention scope should be **PTHREAD\_SCOPE\_SYSTEM**.

**Note:** Caution must be exercised when using the **thread\_setsched** subroutine, since improper use may result in system hangs. See **sys/pri.h** for restrictions on thread priorities.

## Parameters

<i>tid</i>	Specifies the kernel thread ID of the thread whose priority and policy are to be changed.
<i>priority</i>	Specifies the priority to use for this kernel thread. The priority parameter is ignored if the policy is being set to <b>SCHED_OTHER</b> . The priority parameter must have a value in the range 0 to <b>PRI_LOW</b> . <b>PRI_LOW</b> is defined in <b>sys/pri.h</b> . See <b>sys/pri.h</b> for more information on thread priorities.
<i>policy</i>	Specifies the policy to use for this kernel thread. The policy parameter can be one of the following values, which are defined in <b>sys/sched.h</b> : <b>SCHED_OTHER</b> Default operating system scheduling policy. <b>SCHED_FIFO</b> First in-first out scheduling policy. <b>SCHED_FIFO2</b> Allows a thread that sleeps for a relatively short amount of time to be requeued to the head, rather than the tail, of its priority run queue. <b>SCHED_FIFO3</b> Causes threads to be enqueued to the head of their run queues. <b>SCHED_RR</b> Round-robin scheduling policy.

## Return Values

Upon successful completion, the **thread\_setsched** subroutine returns a value of zero. If the **thread\_setsched** subroutine is unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **thread\_setsched** subroutine is unsuccessful if one or more of the following is true:

<b>ESRCH</b>	The kernel thread id <i>tid</i> is invalid.
<b>EINVAL</b>	The policy or priority is invalid.
<b>EPERM</b>	The caller does not have enough privilege to change the policy or priority.

---

## thread\_wait Subroutine

### Purpose

Suspends the thread until it receives a post or times out.

### Library

Standard C library (**libc.a**)

### Syntax

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

```
int thread_wait( timeout)
int timeout;
```

### Description

The **thread\_wait** subroutine allows a thread to wait or block until another thread posts it with the **thread\_post** or the **thread\_post\_many** subroutine or until the time limit specified by the *timeout* value expires. It returns immediately if there is a pending post for this thread or if a *timeout* value of 0 is specified.

If the event for which the thread is waiting and for which it will be posted will occur only in the future, the **thread\_wait** subroutine may be called with a *timeout* value of **0** to clear any pending posts.

The **thread\_wait** and the **thread\_post** subroutine can be used by applications to implement a fast IPC mechanism between threads in different processes.

### Parameters

*timeout*

Specifies the maximum length of time, in milliseconds, to wait for a posting. If the *timeout* parameter value is **-1**, the **thread\_wait** subroutine does not return until a posting actually occurs. If the value of the *timeout* parameter is **0**, the **thread\_wait** subroutine does not wait for a post to occur but returns immediately, even if there are no pending posts. For a non-privileged user, the minimum *timeout* value is 10 msec and any value less than that is automatically increased to 10 msec.

### Return Values

On successful completion, the **thread\_wait** subroutine returns a value of **0**. The **thread\_wait** subroutine completes successfully if there was a pending post or if the calling thread was posted before the time limit specified by the *timeout* parameter expires.

A return value of **THREAD\_WAIT\_TIMEDOUT** indicates that the **thread\_wait** subroutine timed out.

If unsuccessful, a value of **-1** is returned and the global variable **errno** is set to indicate the error.

### Error Codes

The **thread\_wait** subroutine is unsuccessful when one of the following is true:

**EINTR** This subroutine was terminated by receipt of a signal.  
**ENOMEM** There is not enough memory to allocate a timer

## Related Information

The **thread\_post** (“thread\_post Subroutine” on page 358) subroutine, and **thread\_post\_many** (“thread\_post\_many Subroutine” on page 359) subroutine.

---

## tmpfile Subroutine

### Purpose

Creates a temporary file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdio.h>
FILE *tmpfile ( )
```

### Description

The **tmpfile** subroutine creates a temporary file and opens a corresponding stream. The file is opened for update. The temporary file is automatically deleted when all references (links) to the file have been closed.

The stream refers to a file which has been unlinked. If the process ends in the period between file creation and unlinking, a permanent file may remain.

### Return Values

The **tmpfile** subroutine returns a pointer to the stream of the file that is created if the call is successful. Otherwise, it returns a null pointer and sets the **errno** global variable to indicate the error.

### Error Codes

The **tmpfile** subroutine fails if one of the following occurs:

<b>EINTR</b>	A signal was caught during the <b>tmpfile</b> subroutine.
<b>EMFILE</b>	The number of file descriptors currently open in the calling process is already equal to <b>OPEN_MAX</b> .
<b>ENFILE</b>	The maximum allowable number of files is currently open in the system.
<b>ENOSPEC</b>	The directory or file system which would contain the new file cannot be expanded.

## Related Information

The **fopen**, **freopen**, **fdopen** subroutines, **mktemp** subroutine, **tmpnam** or **tempnam** (“tmpnam or tempnam Subroutine”) subroutine, **unlink** (“unlink Subroutine” on page 423) subroutine.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tmpnam or tempnam Subroutine

### Purpose

Constructs the name for a temporary file.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <stdio.h>
char *tmpnam ( String)
char *String;
```

```
char *tempnam ( Directory, FileXPointer)
const char *Directory, *FileXPointer;
```

## Description

**Attention:** The **tmpnam** and **tempnam** subroutines generate a different file name each time they are called. If called more than 16,384 (**TMP\_MAX**) times by a single process, these subroutines recycle previously used names.

The **tmpnam** and the **tempnam** subroutines generate file names for temporary files. The **tmpnam** subroutine generates a file name using the path name defined as **P\_tmpdir** in the **stdio.h** file.

Files created using the **tmpnam** subroutine reside in a directory intended for temporary use. The file names are unique. The application must create and remove the file.

The **tempnam** subroutine enables you to define the directory. The *Directory* parameter points to the name of the directory in which the file is to be created. If the *Directory* parameter is a null pointer or points to a string that is not a name for a directory, the path prefix defined as **P\_tmpdir** in the **stdio.h** file is used. For an application that has temporary files with initial letter sequences, use the *FileXPointer* parameter to define the sequence. The *FileXPointer* parameter (a null pointer or a string of up to 5 bytes) is used as the beginning of the file name.

Between the time a file name is created and the file is opened, another process can create a file with the same name. Name duplication is unlikely if the other process uses these subroutines or the **mktemp** subroutine, and if the file names are chosen to avoid duplication by other means.

## Parameters

<i>String</i>	Specifies the address of an array of at least the number of bytes specified by <b>L_tmpnam</b> , a constant defined in the <b>stdio.h</b> file.  If the <i>String</i> parameter has a null value, the <b>tmpnam</b> subroutine places its result into an internal static area and returns a pointer to that area. The next call to this subroutine destroys the contents of the area.  If the <i>String</i> parameter's value is not null, the <b>tmpnam</b> subroutine places its results into the specified array and returns the value of the <i>String</i> parameter.
<i>Directory</i>	Points to the path name of the directory in which the file is to be created.  The <b>tempnam</b> subroutine controls the choice of a directory. If the <i>Directory</i> parameter is a null pointer or points to a string that is not a path name for an appropriate directory, the path name defined as <b>P_tmpdir</b> in the <b>stdio.h</b> file is used. If that path name is not accessible, the <b>/tmp</b> directory is used. You can bypass the selection of a path name by providing an environment variable, <b>TMPDIR</b> , in the user's environment. The value of the <b>TMPDIR</b> environment variable is a path name for the desired temporary-file directory.
<i>FileXPointer</i>	A pointer to an initial character sequence with which the file name begins. The <i>FileXPointer</i> parameter value can be a null pointer, or it can point to a string of characters to be used as the first characters of the temporary-file name. The number of characters allowed is file system dependent, but 5 bytes is the maximum allowed.

## Return Values

Upon completion, the **tempnam** subroutine allocates space for the string using the **malloc** subroutine, puts the generated path name in that space, and returns a pointer to the space. Otherwise, it returns a null pointer and sets the **errno** global variable to indicate the error. The pointer returned by **tempnam** may be used in the **free** subroutine when the space is no longer needed.

## Error Codes

The **tempnam** subroutine returns the following error code if unsuccessful:

**ENOMEM** Insufficient storage space is available.

**EINVAL** Indicates an invalid *string* value.

## Related Information

The **fopen**, **freopen**, **fdopen** subroutine, **malloc**, **free**, **realloc**, **calloc**, **mallopt**, **mallinfo**, or **alloca** subroutine, **mktemp** or **mkstemp** subroutine, **openx**, **open**, **creat** subroutine, **tmpfile** (“tmpfile Subroutine” on page 363) subroutine, **unlink** (“unlink Subroutine” on page 423) subroutine.

The **environment** file.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## towctrans Subroutine

### Purpose

Character transliteration.

### Library

Standard library (**libc.a**)

### Syntax

```
#include <wctype.h>
wint_t towctrans (wint_t wc, wctrans_t desc) ;
```

### Description

The **towctrans** function transliterates the wide-character code *wc* using the mapping described by *desc*. The current setting of the **LC\_CTYPE** category should be the same as during the call to **wctrans** that returned the value *desc*. If the value of *desc* is invalid (that is, not obtained by a call to **wctrans** or *desc* is invalidated by a subsequent call to **setlocale** that has affected category **LC\_CTYPE**) the result is implementation-dependent.

### Return Values

If successful, the **towctrans** function returns the mapped value of *wc* using the mapping described by *desc*. Otherwise it returns *wc* unchanged.

## Error Codes

The **towctrans** function may fail if:

**EINVAL** desc contains an invalid transliteration descriptor.

## Related Information

The **towlower** (“towlower Subroutine”) subroutine, **towupper** (“towupper Subroutine” on page 367) subroutine, **wctrans** (“wctrans Subroutine” on page 469) subroutine.

The **wctype.h** file.

---

## towlower Subroutine

### Purpose

Converts an uppercase wide character to a lowercase wide character.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <wchar.h>
```

```
wint_t tolower ( WC)  
wint_t WC;
```

### Description

The **towlower** subroutine converts the uppercase wide character specified by the *WC* parameter into the corresponding lowercase wide character. The **LC\_CTYPE** category affects the behavior of the **towlower** subroutine.

### Parameters

*WC* Specifies the wide character to convert to lowercase.

### Return Values

If the *WC* parameter contains an uppercase wide character that has a corresponding lowercase wide character, that wide character is returned. Otherwise, the *WC* parameter is returned unchanged.

## Related Information

The **iswalnum** subroutine, **iswalpha** subroutine, **iswcntrl** subroutine, **iswctype** subroutine, **iswdigit** subroutine, **iswgraph** subroutine, **iswlower** subroutine, **iswprint** subroutine, **iswpunct** subroutine, **iswspace** subroutine, **iswupper** subroutine, **iswxdigit** subroutine, **setlocale** (“setlocale Subroutine” on page 136) subroutine, **towupper** (“towupper Subroutine” on page 367) subroutine, **wctype** (“wctype or get\_wctype Subroutine” on page 470) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

## toupper Subroutine

### Purpose

Converts a lowercase wide character to an uppercase wide character.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <wchar.h>
```

```
wint_t toupper ( WC)  
wint_t WC;
```

### Description

The **toupper** subroutine converts the lowercase wide character specified by the *WC* parameter into the corresponding uppercase wide character. The **LC\_CTYPE** category affects the behavior of the **toupper** subroutine.

### Parameters

*WC* Specifies the wide character to convert to uppercase.

### Return Values

If the *WC* parameter contains a lowercase wide character that has a corresponding uppercase wide character, that wide character is returned. Otherwise, the *WC* parameter is returned unchanged.

### Related Information

The **iswalnum** subroutine, **iswalph** subroutine, **iswcntrl** subroutine, **iswctype** subroutine, **iswdigit** subroutine, **iswgraph** subroutine, **iswlower** subroutine, **iswprint** subroutine, **iswpunct** subroutine, **iswspace** subroutine, **iswupper** subroutine, **iswxdigit** subroutine, **setlocale** (“setlocale Subroutine” on page 136) subroutine, **towlower** (“towlower Subroutine” on page 366) subroutine, **wctype** (“wctype or get\_wctype Subroutine” on page 470) subroutine.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## t\_rcvreldata Subroutine

### Purpose

Receive an orderly release indication or confirmation containing user data.

## Library

## Syntax

```
#include <xti.h>

int t_rcvreldata(
    int fd,
    struct t_discon *discon)
```

## Description

This function is used to receive an orderly release indication for the incoming direction of data transfer and to retrieve any user data sent with the release. The argument *fd* identifies the local transport endpoint where the connection exists, and **discon** points to a **t\_discon** structure containing the following members:

```
struct netbuf udata;
int reason;
int sequence;
```

After receipt of this indication, the user may not attempt to receive more data via **t\_rcv** or **t\_rcvv** (“t\_rcvv Subroutine” on page 369). Such an attempt will fail with **t\_error** set to [TOUTSTATE]. However, the user may continue to send data over the connection if **t\_sndrel** or **t\_sndreldata** (“t\_sndreldata Subroutine” on page 376) has not been called by the user.

The field *reason* specifies the reason for the disconnection through a protocol-dependent reason code, and **udata** identifies any user data that was sent with the disconnection; the field *sequence* is not used.

If a user does not care if there is incoming data and does not need to know the value of *reason*, **discon** may be a null pointer, and any user data associated with the disconnection will be discarded.

If **discon->udata.maxlen** is greater than zero and less than the length of the value, **t\_rcvreldata** fails with **t\_errno** set to [TBUFOVFLW].

This function is an optional service of the transport provider, only supported by providers of service type T\_COTS\_ORD. The flag T\_ORDRELDATA in the *info->flag* field returned by **t\_open** or **t\_getinfo** indicates that the provider supports orderly release user data; when the flag is not set, this function behaves as **t\_rcvrel** and no user data is returned.

This function may not be available on all systems.

Parameters	Before call	After call
fd	x	/
discon->	udata.maxlen	x
discon->	udata.len	/
discon->	udata.buf	?
discon->	reason	/
discon->	sequence	/

## Valid States

T\_DATAXFER, T\_OUTREL

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and **t\_errno** is set to indicate an error.

## Error Codes

On failure, the `t_errno` subroutine is set to one of the following:

### TBADF

The specified file descriptor does not refer to a transport endpoint.

### TBUFOVFLW

The number of bytes allocated for incoming data (`maxlen`) is greater than 0 but not sufficient to store the data, and the disconnection information to be returned in `discon` will be discarded. The provider state, as seen by the user, will be changed as if the data was successfully retrieved.

### TLOOK

An asynchronous event has occurred on this transport endpoint and requires immediate attention.

### TNOREL

No orderly release indication currently exists on the specified transport endpoint.

### TNOTSUPPORT

Orderly release is not supported by the underlying transport provider.

### TOUTSTATE

The communications endpoint referenced by `fd` is not in one of the states in which a call to this function is valid.

### TPROTO

This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (`t_errno`).

### TSYSERR

A system error has occurred during execution of this function.

## Related Information

The `t_getinfo`, `t_open`, `t_sndreldata` (“`t_sndreldata` Subroutine” on page 376), `t_rcvrel`, `t_sndrel` subroutines.

---

## t\_rcvv Subroutine

### Purpose

Receive data or expedited data sent over a connection and put the data into one or more non-contiguous buffers.

### Library

`libxti.*`

### Syntax

```
#include <xti.h>
int t_rcvv (int fd, struct t_iovec *iov, unsigned int iovcount, int *flags) ;
```

### Description

This function receives either normal or expedited data. The argument `fd` identifies the local transport endpoint through which data will arrive, `iov` points to an array of buffer address/buffer size pairs (`iov_base`, `iov_len`). The `t_rcvv` function receives data into the buffers specified by `iov[0].iov_base`, `iov[1].iov_base`, through `iov[iovcount-1].iov_base`, always filling one buffer before proceeding to the next.

**Note:** The limit on the total number of bytes available in all buffers passed (that is, `iov(0).iov_len + . . . + iov(iovcount-1).iov_len`) may be constrained by implementation limits. If no other constraint applies,

it will be limited by [INT\_MAX]. In practice, the availability of memory to an application is likely to impose a lower limit on the amount of data that can be sent or received using scatter/gather functions.

The argument *iovcount* contains the number of buffers which is limited to T\_IOV\_MAX (an implementation-defined value of at least 16). If the limit is exceeded, the function will fail with [TBADDDATA].

The argument *flags* may be set on return from **t\_rcvv** and specifies optional flags as described below.

By default, **t\_rcvv** operates in synchronous mode and will wait for data to arrive if none is currently available. However, if O\_NONBLOCK is set (via **t\_open** or **fcntl**, **t\_rcvv** will execute in asynchronous mode and will fail if no data is available (see [TNODATA] below).

On return from the call, if T\_MORE is set in flags, this indicates that there is more data, and the current transport service data unit (TSDU) or expedited transport service data unit (ETSDU) must be received in multiple **t\_rcvv** or **t\_rcv** calls. In the asynchronous mode, or under unusual conditions (for example, the arrival of a signal or T\_EXDATA event), the T\_MORE flag may be set on return from the **t\_rcvv** call even when the number of bytes received is less than the total size of all the receive buffers. Each **t\_rcvv** with the T\_MORE flag set indicates that another **t\_rcvv** must follow to get more data for the current TSDU. The end of the TSDU is identified by the return of a **t\_rcvv** call with the T\_MORE flag not set. If the transport provider does not support the concept of a TSDU as indicated in the info argument on return from **t\_open** or **ort\_getinfo**, the T\_MORE flag is not meaningful and should be ignored. If the amount of buffer space passed in *iov* is greater than zero on the call to **t\_rcvv**, then **t\_rcvv** will return 0 only if the end of a TSDU is being returned to the user.

On return, the data is expedited if T\_EXPEDITED is set in flags. If T\_MORE is also set, it indicates that the number of expedited bytes exceeded *nbytes*, a signal has interrupted the call, or that an entire ETSDU was not available (only for transport protocols that support fragmentation of ETSDUs). The rest of the ETSDU will be returned by subsequent calls to **t\_rcvv** which will return with T\_EXPEDITED set in flags. The end of the ETSDU is identified by the return of a **t\_rcvv** call with T\_EXPEDITED set and T\_MORE cleared. If the entire ETSDU is not available it is possible for normal data fragments to be returned between the initial and final fragments of an ETSDU.

If a signal arrives, **t\_rcvv** returns, giving the user any data currently available. If no data is available, **t\_rcvv** returns -1, sets **t\_errno** to [TSYSERR] and **errno** to [EINTR]. If some data is available, **t\_rcvv** returns the number of bytes received and T\_MORE is set in flags.

In synchronous mode, the only way for the user to be notified of the arrival of normal or expedited data is to issue this function or check for the T\_DATA or T\_EXDATA events using the **t\_look** function. Additionally, the process can arrange to be notified via the EM interface.

Parameters	Before call	After call
<i>fd</i>	X	/
<i>iov</i>	X/	
<i>iovcount</i>	X	/
<i>iov</i> [0]. <i>iov_base</i>	X(/)	=(X)
<i>iov</i> [0]. <i>iov_len</i>	X	=
...		
<i>iov</i> [ <i>iovcount</i> -1]. <i>iov_base</i>	X(/)	=(X)
<i>iov</i> [ <i>iovcount</i> -1]. <i>iov_len</i>	X	=

## Return Values

On successful completion, **t\_rcvv** returns the number of bytes received. Otherwise, it returns -1 on failure and **t\_errno** is set to indicate the error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

<b>TBADDATA</b>	iovcount is greater than T_IOV_MAX.
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TLOOK</b>	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
<b>TNODATA</b>	O_NONBLOCK was set, but no data is currently available from the transport provider.
<b>TNOTSUPPORT</b>	This function is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The communications endpoint referenced by fd is not in one of the states in which a call to this function is valid.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <b>t_errno</b> ).
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **fcntl** subroutine, **t\_getinfo** subroutine, **t\_look** subroutine, **t\_open** subroutine, **t\_rcv** subroutine, **t\_snd** subroutine, and **t\_sndv** (“t\_sndv Subroutine” on page 373) subroutine.

---

## t\_rcvvdata Subroutine

### Purpose

Receive a data unit into one or more noncontiguous buffers.

### Library

Standard library (**libxti.a**)

### Syntax

```
#include <xti.h>
int t_rcvvdata (
    int fd,      struct t_unitdata *unitdata,      struct t_iovec *iov,      unsigned int iovcount,      int *flags)
```

### Description

This function is used in connectionless mode to receive a data unit from another transport user. The argument **fd** identifies the local transport endpoint through which data will be received, **unitdata** holds information associated with the received data unit, **iovcount** contains the number of non-contiguous udata buffers which is limited to T\_IOV\_MAX (an implementation-defined value of at least 16), and **flags** is set on return to indicate that the complete data unit was not received. If the limit on **iovcount** is exceeded, the function fails with [TBADDATA]. The argument **unitdata** points to a **t\_unitdata** structure containing the following members:

```
    struct netbuf addr;
    struct netbuf opt;
    struct netbuf udata;
```

The **maxlen** field of **addr** and **opt** must be set before calling this function to indicate the maximum size of the buffer for each. The **udata** field of **t\_unitdata** is not used. The **iov\_len** and **iov\_base** fields of **iov[0]**

through **iov[iovcount-1]** must be set before calling **t\_rcvvudata** to define the buffer where the userdata will be placed. If the maxlen field of **addr** or **opt** is set to zero then no information is returned in the **buf** field for this parameter.

On return from this call, **addr** specifies the protocol address of the sending user, **opt** identifies options that were associated with this data unit, and **iov[0].iov\_base** through **iov[iovcount-1].iov\_base** contains the user data that was received. The return value of **t\_rcvvudata** is the number of bytes of user data given to the user.

**Note:** The limit on the total number of bytes available in all buffers passed (that is, **iov(0).iov\_len + . . . + iov(iovcount-1).iov\_len**) may be constrained by implementation limits. If no other constraint applies, it will be limited by [INT\_MAX]. In practice, the availability of memory to an application is likely to impose a lower limit on the amount of data that can be sent or received using scatter/gather functions.

By default, **t\_rcvvudata** operates in synchronous mode and waits for a data unit to arrive if none is currently available. However, if O\_NONBLOCK is set (via **t\_open** or **fcntl**), **t\_rcvvudata** executes in asynchronous mode and fails if no data units are available. If the buffers defined in the **iov[]** array are not large enough to hold the current data unit, the buffers will be filled and T\_MORE will be set in flags on return to indicate that another **t\_rcvvudata** should be called to retrieve the rest of the data unit. Subsequent calls to **t\_rcvvudata** will return zero for the length of the address and options, until the full data unit has been received.

Parameters	Before call	After call
fd	X	/
unitdata->addr.maxlen	X	=
unitdata->addr.len	/	X
unitdata->addr.buf	?(/)	=(/)
unitdata->opt.maxlen	X	=
unitdata->opt.len	/	X
unitdata->opt.buf	?(/)	=(?)
unitdata->udata.maxlen	/	=
unitdata->udata.len	/	=
unitdata->udata.buf	/	=
iov[0].iov_base	X	=(X)
iov[0].iov_len	X	=
. . . .		
iov[iovcount-1].iov_base	X(/)	=(X)
iov[iovcount-1].iov_len	X	=
iovcoun	X	/
flags	/	/

## Return Values

On successful completion, **t\_rcvvudata** returns the number of bytes received. Otherwise, it returns -1 on failure and **t\_errno** is set to indicate the error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

<b>TBADDATA</b>	<b>iovcount</b> is greater than <b>T_IOV_MAX</b> .
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TBUFOVFLW</b>	The number of bytes allocated for the incoming protocol address or options ( <b>maxlen</b> ) is greater than 0 but not sufficient to store the information. The unit data information to be returned in <b>unitdata</b> will be discarded.
<b>TLOOK</b>	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
<b>TNODATA</b>	<b>O_NONBLOCK</b> was set, but no data units are currently available from the transport provider.
<b>TNOTSUPPORT</b>	This function is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The communications endpoint referenced by <b>fd</b> is not in one of the states in which a call to this function is valid.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <b>t_errno</b> ).
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **fcntl** subroutine, **t\_alloc** subroutine, **t\_open** subroutine, **t\_rcvudata** subroutine, **t\_rcvuderr** subroutine, **t\_sndudata** subroutine, **t\_sndvudata** (“**t\_sndvudata** Subroutine” on page 377) subroutine.

---

## t\_sndv Subroutine

### Purpose

Send data or expedited data, from one or more non-contiguous buffers, on a connection.

### Library

Standard library (**libxti.a**)

### Syntax

```
#include <xti.h>
int t_sndv (int fd, const struct t_iovec *iov, unsigned int iovcount, int flags)
```

### Description

Parameters	Before call	After call
fd	X	/
iovec	X	/
iovcount	X	/
iov[0].iov_base	X(X)	/
iov[0].iov_len	X	/
...		
iov[iovcount-1].iov_base	X(X)	/
iov[iovcount-1].iov_len	X	=
flags	X	/

This function is used to send either normal or expedited data. The argument **fd** identifies the local transport endpoint over which data should be sent, **iov** points to an array of buffer address/buffer length

pairs. **t\_sndv** sends data contained in buffers **iov[0]**, **iov[1]**, through **iov[iovcount-1]**. **iovcount** contains the number of non-contiguous data buffers which is limited to **T\_IOV\_MAX** (an implementation-defined value of at least 16). If the limit is exceeded, the function fails with **[TBADDDATA]**.

**Note:** The limit on the total number of bytes available in all buffers passed (that is: **iov(0).iov\_len + . . . + iov(iovcount-1).iov\_len**) may be constrained by implementation limits. If no other constraint applies, it will be limited by **[INT\_MAX]**. In practice, the availability of memory to an application is likely to impose a lower limit on the amount of data that can be sent or received using scatter/gather functions.

The argument **flags** specifies any optional flags described below:

#### **T\_EXPEDITED**

If set in **flags**, the data will be sent as expedited data and will be subject to the interpretations of the transport provider.

#### **T\_MORE**

If set in **flags**, this indicates to the transport provider that the transport service data unit (TSDU) (or expedited transport service data unit ETSDU) is being sent through multiple **t\_sndv** calls. Each **t\_sndv** with the **T\_MORE** flag set indicates that another **t\_sndv** (or **t\_snd**) will follow with more data for the current TSDU (or ETSDU).

The end of the TSDU (or ETSDU) is identified by a **t\_sndv** call with the **T\_MORE** flag not set. Use of **T\_MORE** enables a user to break up large logical data units without losing the boundaries of those units at the other end of the connection. The flag implies nothing about how the data is packaged for transfer below the transport interface. If the transport provider does not support the concept of a TSDU as indicated in the **info** argument on return from **t\_open ort\_getinfo**, the **T\_MORE** flag is not meaningful and will be ignored if set.

The sending of a zero-length fragment of a TSDU or ETSDU is only permitted where this is used to indicate the end of a TSDU or ETSDU, that is, when the **T\_MORE** flag is not set. Some transport providers also forbid zero-length TSDUs and ETSDUs. See Appendix A for a fuller explanation.

If set in **flags**, requests that the provider transmit all data that it has accumulated but not sent. The request is a local action on the provider and does not affect any similarly named protocol flag (for example, the TCP PUSH flag). This effect of setting this flag is protocol-dependent, and it may be ignored entirely by transport providers which do not support the use of this feature.

**Note:** The communications provider is free to collect data in a send buffer until it accumulates a sufficient amount for transmission.

By default, **t\_sndv** operates in synchronous mode and may wait if flow control restrictions prevent the data from being accepted by the local transport provider at the time the call is made. However, if **O\_NONBLOCK** is set (via **t\_open** or **fcntl**), **t\_sndv** executes in asynchronous mode, and will fail immediately if there are flow control restrictions. The process can arrange to be informed when the flow control restrictions are cleared via either **t\_look** or the EM interface.

On successful completion, **t\_sndv** returns the number of bytes accepted by the transport provider. Normally this will equal the total number of bytes to be sent, that is,

$(\text{iov}[0].\text{iov\_len} + \dots + \text{iov}[\text{iovcount}-1].\text{iov\_len})$

However, the interface is constrained to send at most **INT\_MAX** bytes in a single send. When **t\_sndv** has submitted **INT\_MAX** (or lower constrained value, see the note above) bytes to the provider for a single call, this value is returned to the user. However, if **O\_NONBLOCK** is set or the function is interrupted by a signal, it is possible that only part of the data has actually been accepted by the communications provider.

In this case, **t\_sndv** returns a value that is less than the value of **nbytes**. If **t\_sndv** is interrupted by a signal before it could transfer data to the communications provider, it returns -1 with **t\_errno** set to [TSYSERR] and **errno** set to [EINTR].

If the number of bytes of data in the **iov** array is zero and sending of zero octets is not supported by the underlying transport service, **t\_sndv** returns -1 with **t\_errno** set to [TBADDDATA].

The size of each TSDU or ETSDU must not exceed the limits of the transport provider as specified by the current values in the TSDU or ETSDU fields in the **info** argument returned by **t\_getinfo**.

The error [TLOOK] is returned for asynchronous events. It is required only for an incoming disconnect event but may be returned for other events.

## Return Values

On successful completion, **t\_sndv** returns the number of bytes accepted by the transport provider. Otherwise, -1 is returned on failure and **t\_errno** is set to indicate the error.

### Notes:

1. In synchronous mode, if more than INT\_MAX bytes of data are passed in the **iov** array, only the first INT\_MAX bytes will be passed to the provider.
2. If the number of bytes accepted by the communications provider is less than the number of bytes requested, this may either indicate that O\_NONBLOCK is set and the communications provider is blocked due to flow control, or that O\_NONBLOCK is clear and the function was interrupted by a signal.

## Error Codes

On failure, **t\_errno** is set to one of the following:

**TBADDDATA** Illegal amount of data:

- A single send was attempted specifying a TSDU (ETSDU) or fragment TSDU (ETSDU) greater than that specified by the current values of the TSDU or ETSDU fields in the **info** argument.
- A send of a zero byte TSDU (ETSDU) or zero byte fragment of a TSDU (ETSDU) is not supported by the provider.
- Multiple sends were attempted resulting in a TSDU (ETSDU) larger than that specified by the current value of the TSDU or ETSDU fields in the **info** argument the ability of an XTI implementation to detect such an error case is implementation-dependent (see CAVEATS, below).
- **iovcnt** is greater than T\_IOV\_MAX.

<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TBADFLAG</b>	An invalid flag was specified.
<b>TFLOW</b>	O_NONBLOCK was set, but the flow control mechanism prevented the transport provider from accepting any data at this time.
<b>TLOOK</b>	An asynchronous event has occurred on this transport endpoint.
<b>TNOTSUPPORT</b>	This function is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The communications endpoint referenced by <b>fd</b> is not in one of the states in which a call to this function is valid.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error ( <b>t_errno</b> ).
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **t\_getinfo** subroutine, **t\_open** subroutine, **t\_rcvv** (“t\_rcvv Subroutine” on page 369) subroutine, **t\_rcv** subroutine, **t\_snd** subroutine.

---

## t\_sndreldata Subroutine

### Purpose

Initiate/respond to an orderly release with user data.

### Library

### Syntax

```
#include <xti.h>
```

```
int t_sndreldata(int fd, struct t_discon *discon)
```

### Description

This function is used to initiate an orderly release of the outgoing direction of data transfer and to send user data with the release. The argument *fd* identifies the local transport endpoint where the connection exists, and **discon** points to a **t\_discon** structure containing the following members:

```
struct netbuf udata;  
int reason;  
int sequence;
```

After calling **t\_sndreldata**, the user may not send any more data over the connection. However, a user may continue to receive data if an orderly release indication has not been received.

The field **reason** specifies the reason for the disconnection through a protocol-dependent **reason code**, and **udata** identifies any user data that is sent with the disconnection; the field **sequence** is not used.

The **udata** structure specifies the user data to be sent to the remote user. The amount of user data must not exceed the limits supported by the transport provider, as returned in the **discon** field of the *info* argument of **t\_open** or **t\_getinfo**. If the **len** field of **udata** is zero or if the provider did not return T\_ORDRELDATA in the **t\_open** flags, no data will be sent to the remote user.

If a user does not wish to send data and reason code to the remote user, the value of **discon** may be a null pointer.

This function is an optional service of the transport provider, only supported by providers of service type T\_COTS\_ORD. The flag T\_ORDRELDATA in the **info->flag** field returned by **t\_open** or **t\_getinfo** indicates that the provider supports orderly release user data; when the flag is not set, this function behaves as **t\_rcvrel** and no user data is returned.

This function may not be available on all systems.

Parameters	Before call	After call
fd	x	/
discon->	udata.maxlen	/
discon->	udata.len	x
discon->	udata.buf	?(?)
discon->	reason	?
discon->	sequence	/

### Valid States

T\_DATAXFER, T\_INREL

## Error Codes

On failure, **t\_errno** is set to one of the following:

### [TBADDDATA]

The amount of user data specified was not within the bounds allowed by the transport provider, or user data was supplied and the provider did not return T\_ORDRELDATA in the **t\_open** flags.

### [TBADF]

The specified file descriptor does not refer to a transport endpoint.

### [TFLOW]

O\_NONBLOCK was set, but the flow control mechanism prevented the transport provider from accepting the function at this time.

### [TLOOK]

An asynchronous event has occurred on this transport endpoint and requires immediate attention.

### [TNOTSUPPORT]

Orderly release is not supported by the underlying transport provider.

### [TOUTSTATE]

The communications endpoint referenced by **fd** is not in one of the states in which a call to this function is valid.

### [TPROTO]

This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (**t\_errno**).

### [TSYSERR]

A system error has occurred during execution of this function.

## Return Value

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and **t\_errno** is set to indicate an error.

## Related Information

The **t\_getinfo**, **t\_open**, **t\_rcvreldata** (“t\_rcvreldata Subroutine” on page 367), **t\_rcvrel**, and **t\_sndrel** subroutines.

---

## t\_sndvudata Subroutine

### Purpose

Send a data unit from one or more noncontiguous buffers.

### Library

### Syntax

```
#include <xti.h>
```

```
int t_sndvudata(  
    int fd,  
    struct t_unitdata *unitdata,  
    struct t_iovec *iov,  
    unsigned int iovcount)
```

## Description

This function is used in connectionless mode to send a data unit to another transport user. The argument *fd* identifies the local transport endpoint through which data will be sent, **iovcount** contains the number of non-contiguous udata buffers and is limited to an implementation-defined value given by T\_IOV\_MAX, which is at least 16, and **unitdata** points to a **t\_unitdata** structure containing the following members:

```
struct netbuf addr;  
struct netbuf opt;  
struct netbuf udata;
```

If the limit on **iovcount** is exceeded, the function fails with [TBADDDATA].

In **unitdata**, **addr** specifies the protocol address of the destination user, and **opt** identifies options that the user wants associated with this request. The *udata* field is not used. The user may choose not to specify what protocol options are associated with the transfer by setting the *len* field of **opt** to zero. In this case, the provider may use default options.

The data to be sent is identified by **iov[0]** through **iov[iovcount-1]**.

The limit on the total number of bytes available in all buffers passed (that is:

```
iov(0).iov_len + . . . + iov(iovcount-1).iov_len )
```

may be constrained by implementation limits. If no other constraint applies, it will be limited by [INT\_MAX]. In practice, the availability of memory to an application is likely to impose a lower limit on the amount of data that can be sent or received using scatter/gather functions.

By default, **t\_sndvudata** operates in synchronous mode and may wait if flow control restrictions prevent the data from being accepted by the local transport provider at the time the call is made. However, if O\_NONBLOCK is set (via **t\_open** or **fcntl**), **t\_sndvudata** executes in asynchronous mode and will fail under such conditions. The process can arrange to be notified of the clearance of a flow control restriction via either **t\_look** or the EM interface.

If the amount of data specified in **iov[0]** through **iov[iovcount-1]** exceeds the TSDU size as returned in the *tsdu* field of the *info* argument of **t\_open** or **t\_getinfo**, or is zero and sending of zero octets is not supported by the underlying transport service, a [TBADDDATA] error is generated. If **t\_sndvudata** is called before the destination user has activated its transport endpoint (see **t\_bind**), the data unit may be discarded.

If it is not possible for the transport provider to immediately detect the conditions that cause the errors [TBADDADDR] and [TBADDOPT], these errors will alternatively be returned by **t\_rcvuderr**. An application must therefore be prepared to receive these errors in both of these ways.

Parameters	Before call	After call
<i>fd</i>	x	/
<i>unitdata-&gt;addr.maxlen</i>	addr.maxlen	/
<i>unitdata-&gt;addr.len</i>	addr.len	x
<i>unitdata-&gt;addr.buf</i>	addr.buf	x(x)
<i>unitdata-&gt;opt.maxlen</i>	opt.maxlen	/
<i>unitdata-&gt;opt.len</i>	opt.len	x
<i>unitdata-&gt;opt.buf</i>	opt.buf	?(?)
<i>unitdata-&gt;udata.maxlen</i>	udata.maxlen	/
<i>unitdata-&gt;udata.len</i>	udata.len	/
<i>unitdata-&gt;udata.buf</i>	udata.buf	/
<i>iov[0].iov_base</i>	x(x)	=(=)
<i>left&gt;iov[0].iov_len</i>	x	=
. . . .		

## Parameters

iov[iovcnt-1].iov\_base  
iov[iovcnt-1].iov\_len  
iovcnt

## Before call

x(x)  
x  
x

## After call

=(=)  
=  
/

## Valid States

T\_IDLE

## Error Codes

On failure, **t\_errno** is set to one of the following:

### [TBADADDR]

The specified protocol address was in an incorrect format or contained illegal information.

### [TBADDATA]

Illegal amount of data.

- A single send was attempted specifying a TSDU greater than that specified in the *info* argument, or a send of a zero byte TSDU is not supported by the provider.
- **iovcnt** is greater than T\_IOV\_MAX.

### [TBADF]

The specified file descriptor does not refer to a transport endpoint.

### [TBADOPT]

The specified options were in an incorrect format or contained illegal information.

### [TFLOW]

O\_NONBLOCK was set, but the flow control mechanism prevented the transport provider from accepting any data at this time.

### [TLOOK]

An asynchronous event has occurred on this transport endpoint.

### [TNOTSUPPORT]

This function is not supported by the underlying transport provider.

### [TOUTSTATE]

The communications endpoint referenced by **fd** is not in one of the states in which a call to this function is valid.

### [TPROTO]

This error indicates that a communication problem has been detected between XTI and the transport provider for which there is no other suitable XTI error (**t\_errno**).

### [TSYSERR]

A system error has occurred during execution of this function.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and **t\_errno** is set to indicate an error.

## Related Information

The **fcntl**, **t\_alloc**, **t\_open**, **t\_rcvudata**, **t\_rcvvudata** (“t\_rcvvudata Subroutine” on page 371), **t\_rcvuderr**, **t\_sndudata** subroutines.

---

## t\_sysconf Subroutine

### Purpose

Get configurable XTI variables.

### Library

Standard library (**libxti.a**)

### Syntax

```
#include <xti.h>
```

```
int t_sysconf ( int name)
```

## Description

Parameters	Before call	After call
name	X	/

The **t\_sysconf** function provides a method for the application to determine the current value of configurable and implementation-dependent XTI limits or options.

The **name** argument represents the XTI system variable to be queried. The following table lists the minimal set of XTI system variables from **xti.h** that can be returned by **t\_sysconf**, and the symbolic constants, defined in **xti.h** that are the corresponding values used for **name**.

Variable	Value of Name
T_IOV_MAX	_SC_T_IOV_MAX

## Return Values

If **name** is valid, **t\_sysconf** returns the value of the requested limit/option (which might be -1) and leaves **t\_errno** unchanged. Otherwise, a value of -1 is returned and **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to the following:

**TBADFLAG** **name** has an invalid value.

## Related Information

The **t\_rcvv** (“t\_rcvv Subroutine” on page 369) subroutine, **t\_rcvvudata** (“t\_rcvvudata Subroutine” on page 371) subroutine, **t\_sndv** (“t\_sndv Subroutine” on page 373) subroutine, **t\_sndvudata** (“t\_sndvudata Subroutine” on page 377) subroutine.

---

## trc\_close Subroutine

### Purpose

Closes and frees a trace log object.

### Library

**libtrace.a**

### Syntax

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

```
int trc_close (handle)  
trc_log_handle_t handle;
```

## Description

The **trc\_close** subroutine closes a trace log object. The object must have been opened with the **trc\_open** subroutine. If the **TRC\_RETAIN\_HANDLE** type was specified at open time, the **trc\_close** subroutine must be called after a call to the **trc\_open** subroutine, regardless of whether the open succeeded or not.

## Parameters

*handle* Contains the handle returned from a successful call to the **trc\_open** subroutine.

## Return Values

Upon successful completion, the **trc\_close** subroutine returns a 0.

## Error Codes

Upon error, the **trc\_close** subroutine sets the **errno** global variable and returns the error from the **fclose** subroutine. In addition, **EINVAL** is returned if *handle* contains an invalid **trc\_log\_handle\_t** object.

## Related Information

“trc\_open Subroutine” on page 392, “trc\_read Subroutine” on page 395, “trc\_loginfo Subroutine” on page 390, “trc\_find\_first, trc\_find\_next, and trc\_compare Subroutine”, “trc\_seek and trc\_tell Subroutine” on page 398, “trc\_libcntl Subroutine” on page 388, “trc\_strerror Subroutine” on page 399, “trc\_perror Subroutine” on page 394, “trcstart Subroutine” on page 405, “trcon Subroutine” on page 404, “trcoff Subroutine” on page 404 and “trcstop Subroutine” on page 406.

The trace daemon in *AIX 5L Version 5.2 Commands Reference, Volume 5*.

The **trcrpt**, **trcstop**, and **trcupdate** commands in *AIX 5L Version 5.2 Commands Reference, Volume 5*.

---

## trc\_find\_first, trc\_find\_next, and trc\_compare Subroutine

### Purpose

Finds the first, or next, occurrence of the argument, or compares the current entry with the argument.

### Library

**libtrace.a**

### Syntax

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

```
int trc_find_first (handle, argp, ret)
trc_log_handle_t handle;
trc_logsearch_t *argp;
trc_read_t *ret;
```

```
int trc_find_next (handle, argp, ret)
trc_log_handle_t handle;
trc_logsearch_t *argp;
trc_read_t *ret;
```

```
int trc_compare (handle, argp)
trc_log_handle_t handle;
trc_logsearch_t *argp;
```

### Description

The **trc\_find\_first** subroutine finds the first occurrence of the trace log entry matching the argument pointed to by the *argp* parameter. The **trc\_find\_next** subroutine finds the next occurrence of the argument starting from the current position in the log object. If the search argument pointer, *argp*, is NULL, the argument from the previous search is used. Both the **trc\_find\_first** and **trc\_find\_next** subroutines return the item found.

The **trc\_compare** subroutine is used to check the current entry against the argument. No data is read. It is useful when implementing exit criteria, where you need to find entries according to some criteria, but then check for an exit criteria which is not part of the normal search.

## Parameters

<i>handle</i>	Contains the handle returned from a successful call to the <b>trc_open</b> subroutine.
<i>argp</i>	Points to the argument list as defined in the <b>/usr/include/sys/libtrace.a</b> file. Arguments may be chained together to perform complex searches.
<i>ret</i>	Points to the <b>trc_read_t</b> structure to be returned. The <b>trc_free</b> subroutine should be used to free data referenced from the <b>trc_read_t</b> data type, unless <b>TRC_LOGLIVE</b> was specified at open time.

The search argument consists of three parts, the operator, **tls\_op**, and the left and right sides.

The operator values can be easily identified, because they have the form **TLS\_OP\_...** Operators are split into two categories, leaf and compound operators. Leaf operators are operators that compare the field on the left with the value on the right. Compound operators are used to compare two expressions, (for example) to combined expressions.

Leaf operations may be performed using numeric or string data. If performed on string data, the **strcmp** **libc** string compare function is used to do the comparison for all operators except **TLS\_OP\_SUBSTR**. The valid leaf operators are:

<b>TLS_OP_EQUAL</b>	Exactly equal
<b>TLS_OP_NE</b>	Not equal
<b>TLS_OP_LT</b>	Less than
<b>TLS_OP_LE</b>	Less than or equal
<b>TLS_OP_GT</b>	Greater than
<b>TLS_OP_GE</b>	Greater than or equal
<b>TLS_OP_SUBSTR</b>	The string on the left contains the string on the right.

The compound operators are:

<b>TLS_OP_AND</b>	The logical AND of the results of the left and right expressions.
<b>TLS_OP_OR</b>	The logical OR of the results of the left and right expressions.
<b>TLS_OP_XOR</b>	The exclusive or of the results of the left and right expressions.
<b>TLS_OP_NOT</b>	The negation of the argument referenced by <b>tls_left</b> .

The left and right sides of the expression are defined as follows:

<b>tls_left</b> and <b>tls_right</b>	These are used when the operator requires the left and right sides to be an expression, (for example) when it is a compound operator. <b>tls_left</b> and <b>tls_right</b> point to other <b>trc_logsearch_t</b> structures.
<b>tls_field</b> and corresponding values	For a leaf operation, <b>tls_field</b> , on the left, specifies the field to be compared. The field names can be identified easily, because they all have the form <b>TLS_MATCH_...</b> The righthand side is a value specified according to the data type of the field on the left.

The following table shows the lefthand field values and their corresponding righthand side data values:

Field	Value	Description
<b>TLS_MATCH_HOOKID</b>	<b>tls_shortvalue</b>	Compare the hookid with a short data item.
<b>TLS_MATCH_HOOK_AND_SUBHOOK</b>	<b>tls_intvalue</b>	Compare the hook and subhook, 28 bits, with the specified integer. Note that the field is of the form 0x0hhhssss, where hhh is the hook id, and ssss is the subhook.
<b>TLS_MATCH_HOOKSET</b>	<b>tls_hooksetvalue</b>	The bit map specifying the hooks to be tested for. This allows you to test for multiple hooks with one search argument. The bit map is manipulated with the <b>trc_hkemptyset</b> , <b>trc_hkfillset</b> , <b>trc_hkaddset</b> , and <b>trc_hkdelset</b> subroutines.
<b>TLS_MATCH_TIME</b>	<b>tls_longvalue</b>	Compare the time value in nanoseconds from the start of the trace.
<b>TLS_MATCH_TID</b>	<b>tls_longvalue</b>	Thread id
<b>TLS_MATCH_PID</b>	<b>tls_longvalue</b>	Process id
<b>TLS_MATCH_RAWOFST</b>	<b>tls_longvalue</b>	Raw file offset
<b>TLS_MATCH_CPUID</b>	<b>tls_intvalue</b>	cpu id
<b>TLS_MATCH_RCPU</b>	<b>tls_intvalue</b>	Remaining cpus in the trace.
<b>TLS_MATCH_FLAGS</b>	<b>tls_intvalue</b>	Compare with <b>trcr_flags</b>
<b>TLS_MATCH_INTR_DEPTH</b>	<b>tls_intvalue</b>	Compare with <b>trchi_intr_depth</b>
<b>TLS_MATCH_PROCNAME</b>	<b>tls_strvalue</b>	Process name
<b>TLS_MATCH_SVCNAME</b>	<b>tls_strvalue</b>	svc name
<b>TLS_MATCH_PRI</b>	<b>tls_intvalue</b>	Dispatch priority
<b>TLS_MATCH_TICKS</b>	<b>tls_longvalue</b>	Match with the number of timer register ticks since the start of the trace.
<b>TLS_MATCH_DATA</b>	<b>tls_strvalue</b>	Compare string with the ascii data, <b>trchi_ascii</b>
<b>TLS_MATCH_FILENAME</b>	<b>tls_strvalue</b>	Compare with <b>trchi_filename</b>
<b>TLS_MATCH_TRCONTIME</b>	<b>tls_longvalue</b>	Compare with <b>trchi_trcontime</b>
<b>TLS_MATCH_TRCOFFTIME</b>	<b>tls_longvalue</b>	Compare with <b>trchi_trcofftime</b>

## Return Values

Upon successful completion, the **trc\_find\_first**, **trc\_find\_next**, and **trc\_compare** subroutines return 0.

## Error Codes

Upon error, the **errno** global variable is set to a value from the **errno.h** file. The **trc\_find\_first**, **trc\_find\_next**, and **trc\_compare** subroutines return either a value from the **errno.h** file, or an error value from the **libtrace.h** file.

**EINVAL**                      The handle is invalid, or the search argument is invalid.

<b>TRCE_EOF</b>	No matching item was found, or no more matching items exist. The <b>errno</b> global variable is set to 0.
<b>TRCE_BADFORMAT</b>	The log object contains badly formatted data. The <b>errno</b> global variable is set to <b>EINVAL</b> .

## Examples

1. Find the SVC hooks, 101 and 104, for program mypgm.

```

{
    int rv;
    trc_loghandle_t h;
    trc_read_t r;
    trc_logsearch_t t1, t2, t3, t4, t5;

    /* Setup the leaf search arguments. */
    t1.tls_op = TLS_OP_EQUAL;
    t1.tls_field = TLS_MATCH_HOOKID;
    t1.tls_shortvalue = 0x101;
    t2.tls_op = TLS_OP_EQUAL;
    t2.tls_field = TLS_MATCH_HOOKID;
    t2.tls_shortvalue = 0x104;
    t3.tls_op = TLS_OP_EQUAL;
    t3.tls_field = TLS_MATCH_PROCNAME;
    t3.tls_strvalue = "mypgm";
    /* Join the items and form a single search tree. */
    t4.tls_op = TLS_OP_AND;
    t4.tls_left = &t1;
    t4.tls_right = &t2;
    t5.tls_op = TLS_OP_AND;
    t5.tls_left = &t4;
    t5.tls_right = &t3;
    /* Open the default trace log object. */
    rv = trc_open("", "", TRC_LOGREAD|TRC_LOGPROC, >h);
    if (rv) {
        trc_perror(h, rv, "open");
        return(rv);
    }
    /* Do the search. */
    rv = trc_find_first(h, &t5, &r);
    if (rv) {
        trc_perror(h, rv, "find test");
        return(rv);
    }
    ...
}

```

Note that subsequent entries matching this search could be returned with the following:

```
rv = trc_find_next(h, NULL, &r);
```

After a find, **trc\_find\_next** can be used to change the search argument without starting the search over. In other words, **trc\_find\_first** always starts from the beginning of the file, while **trc\_find\_next** starts from the current position in the file, but either one can change the search argument.

2. Find the SVC hooks, 101 and 104, for program mypgm. Use a single argument to search for both hook ids.

```

{
    int rv;
    trc_loghandle_t h;
    trc_read_t r;
    trc_logsearch_t t1, t2, t3;
    trc_hookset_t hs;

    /* Setup the hook set. */

```

```

trc_hkemptyset(hs);
(void)trc_hkaddset(hs, 0x101);
(void)trc_hkaddset(hs, 0x104);
/* Setup the leaf search arguments. */
t1.tls_op = TLS_OP_EQUAL;
t1.tls_field = TLS_MATCH_HOOKSET;
t1.tls_hooksetvalue = hs;
t2.tls_op = TLS_OP_EQUAL;
t2.tls_field = TLS_MATCH_PROCNAME;
t2.tls_strvalue = "mypgm";
/* Join the items and form a single search tree. */
t3.tls_op = TLS_OP_AND;
t3.tls_left = &t1
t3.tls_right = &t2
/* Open the default trace log object. */
rv = trc_open("", "", TRC_LOGREAD|TRC_LOGPROC, &h);
if (rv) {
    trc_perror(h, rv, "open");
    return(rv);
}
/* Do the search. */
rv = trc_find_first(h, &t3, &r);
if (rv) {
    trc_perror(h, rv, "find test");
    return(rv);
}
...
}

```

## Related Information

“trc\_open Subroutine” on page 392, “trc\_close Subroutine” on page 380, “trc\_read Subroutine” on page 395, “trc\_loginfo Subroutine” on page 390, “trc\_seek and trc\_tell Subroutine” on page 398, “trc\_libcntl Subroutine” on page 388, “trc\_strerror Subroutine” on page 399, “trc\_perror Subroutine” on page 394, “trcstart Subroutine” on page 405, “trcon Subroutine” on page 404, “trcoff Subroutine” on page 404, “trcstop Subroutine” on page 406, and “trc\_hkemptyset, trc\_hkfillset, trc\_hkaddset, trc\_hkdelset, and trc\_hkisset Subroutine” on page 386.

The trace daemon in *AIX 5L Version 5.2 Commands Reference, Volume 5*.

The trcrpt, trcstop, and trcupdate commands in *AIX 5L Version 5.2 Commands Reference, Volume 5*.

---

## trc\_free Subroutine

### Purpose

Frees memory allocated by the **trc\_read**, **trc\_find**, **trc\_loginfo**, or **trc\_hookname** subroutine.

### Library

**libtrace.a**

### Syntax

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

```
int trc_free (parmp)
void *parmp;
```

## Description

The **trc\_free** subroutine is used to free memory associated with data structures returned by the trace retrieval API. It does not free the storage for the base structure, however, only storage allocated by the API on behalf of the user. The pointer must point to one of the following:

### **trc\_read\_t**

Data returned by the **trc\_read** or **trc\_find** subroutine.

### **trc\_loginfo\_t**

Data returned by the **trc\_loginfo** subroutine.

### **trc\_hookname\_t**

Data returned by the **trc\_hookname** subroutine.

### **trc\_logpos\_t**

A log position object returned by the **trc\_tell** subroutine.

A log handle, **trc\_loghandle\_t**, must be freed using the **trc\_close** subroutine.

For example, `trc_free(&trc_data)`, where `trc_data` is of type **trc\_read\_t**, frees the storage referenced by the **trc\_data** structure, but does not free **trc\_data** since it must be pre-allocated by the user.

## Parameters

*parm* Points to a structure as described above.

## Return Values

Upon successful completion, the **trc\_free** subroutine returns 0.

## Error Codes

**EINVAL** The *parm* parameter points to an unsupported data type.

## Related Information

“trc\_read Subroutine” on page 395, “trc\_loginfo Subroutine” on page 390, “trc\_find\_first, trc\_find\_next, and trc\_compare Subroutine” on page 381, “trc\_hookname Subroutine” on page 387, “trc\_seek and trc\_tell Subroutine” on page 398, “trc\_strerror Subroutine” on page 399, and “trc\_perror Subroutine” on page 394.

---

## trc\_hkemptyset, trc\_hkfillset, trc\_hkaddset, trc\_hkdelset, and trc\_hkisset Subroutine

### Purpose

Manipulates a trace hook set.

### Library

**libtrace.a**

### Syntax

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

```
void trc_hkemptyset(hookset)
trc_hookset_t hookset;
```

```

void trc_hkfillset(hookset)
trc_hookset_t hookset;

int trc_hkaddset(hookset, hook)
trc_hookset_t hookset;
short hook;

int trc_hkdelset(hookset, hook)
trc_hookset_t hookset;
short hook;

int trc_hkisset (hookset, hook)
trc_hookset_t hookset;
short hook

```

## Description

These subroutines manipulate a trace hook set used by the **trc\_find** subroutines. This hook set can be used to search for several trace hooks simultaneously.

## Parameters

*hookset*               References the hook set to be operated on.  
*hook*                    Specifies a hook value in the range 0x000 - 0xffff.

## Return Values

The **trc\_hkaddset**, **trc\_hkdelset**, and **trc\_hkisset** subroutines return **EINVAL** if the hook is out of range (that is, greater than 0xffff).

The **trc\_hkaddset** subroutine returns 0 if the hook wasn't in the set, and -1 if it was already present.

The **trc\_hkdelset** subroutine returns 0 if the hook was in the set, and -1 if it wasn't present.

The **trc\_hkisset** subroutine returns 0 if the hook isn't present, and -1 if it is present.

## Related Information

"trc\_loginfo Subroutine" on page 390 and "trc\_find\_first, trc\_find\_next, and trc\_compare Subroutine" on page 381.

## trc\_hookname Subroutine

### Purpose

Returns one or all hooks and associated names from the template file.

### Library

libtrace.a

### Syntax

```

#include <sys/libtrace.h>

int trc_hookname (handle, hook, hooknamep)
trc_log_handle_t handle;
trc_hookid_t hook;
trc_hookname_t *hooknamep;

```

## Description

The `trc_hookname` subroutine returns one or more hook ids and their associated descriptions. This allows a trace data formatter to provide a hook selection list with some descriptive text for each hook.

## Parameters

<i>handle</i>	Contains a <code>trc_log_handle_t</code> data item returned from a successful call to the <code>trc_open</code> subroutine.
<i>hook</i>	Contains a hook id of the form 0xhhh where hhh is the 3-hex-digit hook id. If the <i>hook</i> parameter is <code>TRC_HOOK_ALL</code> , the names for all hooks in the template file are returned.
<i>hooknamep</i>	Points to a <code>trc_hookname_t</code> structure. The <code>trc_free</code> subroutine should be used to free any data referenced by the <code>trc_hookname_t</code> data item.

```
/* Array element type for hook ids and names. */
typedef struct {
    trc_hookid_t hookid;
    char *hookname;
} trc_hooknm_t;

typedef struct {
    int trchn_magic;           /* Identifier for this data structure. */
    unsigned trchn_nhooks;    /* Number of hooks. */
    trc_hooknm_t *trchn_names; /* Pointer to array of ids and names. */
} trc_hookname_t;
```

## Return Values

Upon successful completion, the `trc_hookname` subroutine returns 0.

## Error Codes

<code>ENOMEM</code>	Not enough memory to satisfy the request.
<code>TRCE_WARN</code>	A formatting error was found in the template file. If <code>TRCE_WARN</code> is returned, the function completed.
<code>TRCE_BADFORMAT</code>	A formatting error was found in the template file. If <code>TRCE_BADFORMAT</code> was returned, the <code>errno</code> global variable is set to <code>EINVAL</code> .

## Related Information

“`trc_open` Subroutine” on page 392, “`trc_loginfo` Subroutine” on page 390, “`trc_free` Subroutine” on page 385, “`trc_strerror` Subroutine” on page 399, and “`trc_perror` Subroutine” on page 394.

---

## `trc_libcntl` Subroutine

### Purpose

Performs trace API control functions.

### Library

`libtrace.a`

## Syntax

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

```
int trc_libcntl (handle, cmd, datap)  
trc_log_handle_t handle;  
int cmd;  
void *datap;
```

## Description

The `trc_libcntl` subroutine provides miscellaneous control functions.

## Parameters

*handle*

Contains the handle returned from a successful call to the `trc_open` subroutine.

*cmd*

This is the control function to be performed. Supported functions are:

### TRC\_CNTL\_ADJLINENO

This allows a trace report program to adjust the `$LINENO` value supplied through the trace templates. Normally, a trace reporting program may assume the `$LINENO` value is calculated based upon the first line of the output, in `trchi_ascii`, being the first line printed for that hook in the report. If this is not the case, such as with the `2line trcrpt` option, the `$LINENO` value must be adjusted.

For `TRC_CNTL_ADJLINENO`, the *datap* parameter must contain a signed long value which is added to `$LINENO`. If the value is negative, `TRC_CNTL_ADJLINENO` will decrement the value.

### TRC\_CNTL\_NAMELIST

This allows the namelist to be specified. The default is `/unix`. It does not initialize the symbols, however, and the `trc_libcntl` subroutine returns `EINVAL` if the symbols are already initialized. If symbols are in the trace stream, specified by `trace -n`, those symbols are used regardless of the namelist specification.

### TRC\_CNTL\_TEXTOFFSET

This offsets each line of text, in the `trchi_ascii` data area, by the number of character positions specified, plus  $(trchi\_indent-1) * 8$ ; If the associated value is 0, each line is only offset by  $(trchi\_indent-1) * 8$ ;

### TRC\_CNTL\_TEXTOFFSET\_SUBSEQUENT

This works exactly like `TRC_CNTL_TEXTOFFSET`, except it offsets all lines except the first line of text. The first line is still offset by  $(trchi\_indent-1) * 8$ ;

### TRC\_CNTL\_PAGESIZE

This specifies the length of a page.

### TRC\_CNTL\_TEXTHEADER

This specifies a header to be output every page, as specified by the `TRC_CNTL_PAGESIZE` command.

*datap*

Specifies the data parameter.

## Return Values

Upon successful completion, the `trc_libcntl` subroutine returns 0.

## Error Codes

**EINVAL**

The *handle* or *cmd* parameter is invalid. **EINVAL** is also returned if the value specified with `TRC_CNTL_ADJLINENO` would cause the `$LINENO` value to be negative.

## Related Information

“trc\_open Subroutine” on page 392, “trc\_close Subroutine” on page 380, “trc\_read Subroutine” on page 395, “trc\_loginfo Subroutine”, “trc\_find\_first, trc\_find\_next, and trc\_compare Subroutine” on page 381, “trcstart Subroutine” on page 405, “trcon Subroutine” on page 404, “trcoff Subroutine” on page 404 and “trcstop Subroutine” on page 406.

The trace daemon in *AIX 5L Version 5.2 Commands Reference, Volume 5*.

The trcrpt, trcstop, and trcupdate commands in *AIX 5L Version 5.2 Commands Reference, Volume 5*.

---

## trc\_loginfo Subroutine

### Purpose

Returns information about a trace log object.

### Library

libtrace.a

### Syntax

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

```
int trc_loginfo (log_object_name, infop)
char *handle;
trc_log_info_t *infop;
```

### Description

The **trc\_loginfo** subroutine returns information about the named trace log object. If the *log\_object\_name* parameter is NULL or an empty string, the **trc\_loginfo** subroutine returns information about the default log object.

### Parameters

<i>log_object_name</i>	Names the trace log object. This is specified as it is for the <b>trc_open</b> subroutine.
<i>infop</i>	Points to an item of type <b>trc_log_info_t</b> where the information will be returned. The <b>trc_log_info_t</b> structure is defined in the <code>/usr/include/sys/libtrace.h</code> file. It contains such fields as the file size, the time the trace was taken, the trace log file magic number, the command used to start the trace, cpus in the machine, number of cpus traced, multicpu trace indicator (-C), and the trace object type as defined in the <b>trcopen</b> subroutine. The <b>trc_free</b> subroutine should be called to free the <b>trc_loginfo_t</b> information, even if the <b>trc_loginfo</b> subroutine returned an error.

The `/usr/include/sys/libtrace.h` file contains the data definitions for the returned data, **\*infop**. The following table contains the data item name, data type, and description for each item returned:

Label	Data Type	Description
trci_magic	int	Structure magic number managed by the library.

Label	Data Type	Description
<b>trci_logmagic</b>	int	The trace log file's magic number, see the <code>/usr/include/sys/trchdr.h</code> file. This identifies the type of log file, and is included mainly for completeness. The pertinent log file information may be gotten from other fields in this structure.
<b>trci_time</b>	time_t	The time the trace was taken.
<b>trci_ipaddr</b>	int	The system's IP address.
<b>trci_uname</b>	struct utsname	uname information.
<b>trci_cmd</b>	char *	The command used to start the trace.
<b>trci_fnames</b>	trci_fname_t*	Log file names array.
<b>trci_mach_cpus</b>	int	Number of cpus in the machine.
<b>trci_traced_cpus</b>	int	Number of traced cpus.
<b>trci_flags</b>	int	Data stream flags.
<b>trci_obj_type</b>	int	Trace object type.
<b>trci_hookids</b>	trc_hookset_t	Binary hook ids map showing the hooks traced. This can be examined with the <b>trc_hkisset</b> subroutine.

The **trci\_flags** field contains bit flags as follows:

**TRCIF\_MULTICPU**

This trace was taken with the **-C** trace option, (for example) it is a multi-cpu trace.

**TRCIF\_64BIT**

This is a 64-bit trace, 32-bit if not set.

**TRCIF\_SEPSEG**

Separate segment buffering was used.

**TRCIF\_CONDTRACE**

Conditional trace by hookid, trace **-j**, **-k**, **-J**, or **-K**.

**TRCIF\_CONDEXCL**

Trace hook exclusion, **-k** or **-K**, was used.

## Return Values

Upon successful completion, the **trc\_loginfo** subroutine returns a 0, and information about the trace log object is placed into the memory pointed to by the *infor* parameter.

## Error Codes

Upon error, the **trc\_loginfo** subroutine returns information identical to that returned by the "trc\_open Subroutine" on page 392.

## Related Information

"trc\_open Subroutine" on page 392, "trc\_close Subroutine" on page 380, "trc\_read Subroutine" on page 395, "trc\_find\_first, trc\_find\_next, and trc\_compare Subroutine" on page 381, "trc\_seek and trc\_tell Subroutine" on page 398, "trc\_libcntl Subroutine" on page 388, "trc\_strerror Subroutine" on page 399, "trc\_perror Subroutine" on page 394, "trcstart Subroutine" on page 405, "trcon Subroutine" on page 404, "trcoff Subroutine" on page 404, "trcstop Subroutine" on page 406, and "trc\_hkemptyset, trc\_hkfillset, trc\_hkaddset, trc\_hkdelset, and trc\_hkisset Subroutine" on page 386.

The trace daemon in *AIX 5L Version 5.2 Commands Reference, Volume 5*.

The trcrpt, trcstop, and trcupdate commands in *AIX 5L Version 5.2 Commands Reference, Volume 5*.

---

## trc\_open Subroutine

### Purpose

Opens a trace log object.

### Library

libtrace.a

### Syntax

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

int trc_open (log_object_name, template_file_name, type, handlep)
char *log_object_name, template_file_name;
int type;
trc_log_handle_t *handlep;
```

### Description

The **trc\_open** subroutine opens a trace log object. A log object may only be opened for reading.

Two object types are supported, raw and processed. As their names imply, a raw object consists of the raw trace data as it was traced. A processed object consists of data as processed by a trace formatting template file such as the **/etc/trcfmt** file.

### Parameters

<i>log_object_name</i>	Specifies the log object to be opened. If this is NULL or an empty string, the default log object, <b>/var/adm/ras/trcfile</b> , is opened. If it is a dash, the input is read from standard input. In this case, the file must be a sequential trace file such as one produced by the <b>trcrpt -r</b> command, the <b>-o</b> trace option, or the <b>trcdead</b> command.  If the file is the base file for a multi-cpu trace, the trace events are merged by the <b>trcrpt</b> command, unless the <b>TRC_NOTEMPLATES</b> option was specified. Also, if the file is a single cpu's trace file, it is treated as a single log file.
<i>template_file_name</i>	This names the template file. The template file is used if the <b>TRC_LOGPROC</b> type is specified. If NULL, <b>/etc/trcfmt</b> (the default template file) is used. The template file specification is ignored if the <b>TRC_NOTEMPLATES</b> option is specified.

*type*

Consists of flag bits OR'd together. One open type and one object type flag must be specified.

The following is the open type flag:

**TRC\_LOGREAD**

Open for reading

The following are the object type flags:

**TRC\_LOGRAW**

Specifies that raw trace data is to be read. This data is defined in Debug and Performance Tracing and in the **/etc/trcfmt** file.

**TRC\_LOGPROC**

This processes a raw trace log file, one produced by the **trace** command, using either the trace templates found in the **/etc/trcfmt** file, or the template file specified by the *template\_file\_name* parameter on the **trc\_open** command.

The following are the modifier type flags:

**TRC\_LOGVERBATIM**

Returns the file data verbatim, exactly as traced. This is how **trcrpt -r** returns data. See also the **TRC\_NOTEMPLATES** modifier.

**TRC\_LIBDEBUG**

Turns on debug mode. This is for IBM customer support use only.

**TRC\_LOGLIVE**

The data returned in the **trc\_read\_t** structure is not a unique copy, it is live data. Such data may only be used until the next retrieval API operation. It is not necessary to call the **trc\_free** subroutine to free such data. The **TRC\_LOGLIVE** modifier is used to improve performance when the data read does not need to be retained.

**TRC\_RETAIN\_HANDLE**

Don't free the handle after an open failure. This allows errors to be processed by the **trc\_perror** or **trc\_strerror** subroutines. The **trc\_close** subroutine must be used to free the file handle.

**TRC\_NOTEMPLATES**

Ignore any template file. This is used with the **TRC\_LOGRAW** object flag to prevent any template processing, such as merging multi-cpu trace files. When used in conjunction with the **TRC\_LOGVERBATIM** flag, it causes the retrieval API to return the same data reported with **trcrpt -r**.

*handlep*

Points to the handle returned from a successful call to the **trc\_open** subroutine.

## Return Values

Upon successful completion, the **trc\_open** subroutine returns a 0 and puts the trace log object handle into the memory pointed to by the *handlep* parameter.

## Error Codes

Upon error, the **trc\_open** subroutine sets the **errno** global variable to a value in the **errno.h** file, and returns either an **errno.h** value, or an error value defined in the **libtrace.h** file.

**EINVAL**

Invalid parameter.

**ENOMEM**

Cannot allocate memory.

**TRCE\_BADFORMAT**

The file is not a valid trace file, and **errno** is set to **EINVAL**.

**TRCE\_WARN**

The template file contains errors. The **errno** global variable is set to **EINVAL** if **TRCE\_TMPLTFORMAT** is returned. If **TRCE\_WARN** is returned, the open succeeded.

<b>TRCE_TMPLTFORMAT</b>	The template file contains errors. The <b>errno</b> global variable is set to <b>EINVAL</b> if <b>TRCE_TMPLTFORMAT</b> is returned. If <b>TRCE_WARN</b> is returned, the open succeeded.
<b>TRCE_TOOMANY</b>	An internal limit is exceeded. The <b>errno</b> global variable is set to <b>ENOMEM</b> in this case.

## Related Information

“trc\_close Subroutine” on page 380, “trc\_read Subroutine” on page 395, “trc\_loginfo Subroutine” on page 390, “trc\_find\_first, trc\_find\_next, and trc\_compare Subroutine” on page 381, “trc\_seek and trc\_tell Subroutine” on page 398, “trc\_libcntl Subroutine” on page 388, “trc\_strerror Subroutine” on page 399, “trc\_perror Subroutine”, “trcstart Subroutine” on page 405, “trcon Subroutine” on page 404, “trcoff Subroutine” on page 404 and “trcstop Subroutine” on page 406.

The trace daemon in *AIX 5L Version 5.2 Commands Reference, Volume 5*.

The trcrpt, trcstop, and trcupdate commands in *AIX 5L Version 5.2 Commands Reference, Volume 5*.

---

## trc\_perror Subroutine

### Purpose

Prints all errors associated with a trace log object.

### Library

libtrace.a

### Syntax

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

```
void trc_perror (handle, rv, str)
void *handle;
int rv;
char *str;
```

### Description

The **trc\_perror** subroutine works like the  **perror** subroutine. If the error in the *rv* parameter is an error from the **errno.h** file, it behaves exactly like the  **perror** subroutine.

If there are multiple errors associated with the handle, the **trc\_perror** subroutine prints all errors associated with the object. If the *str* parameter is NULL, the error’s text is the only text printed. Errors are printed to standard error.

### Parameters

<i>handle</i>	Contains the handle returned from the call to the <b>trc_open</b> subroutine, the <b>trc_logpos_t</b> object returned by the call to the <b>trc_loginfo</b> subroutine, or NULL. If a handle returned by the <b>trc_open</b> subroutine is passed, the <b>trc_open</b> subroutine need not have been successful, and the <b>TRC_RETAIN_HANDLE</b> option must have been used.
<i>rv</i>	The return value from a <b>libtrace</b> subroutine.
<i>str</i>	Used the same as the string passed to the <b> perror</b> subroutine. Errors printed by the <b>trc_perror</b> subroutine are printed as <i>str</i> : error-message.

## Related Information

“trc\_open Subroutine” on page 392, “trc\_read Subroutine”, “trc\_loginfo Subroutine” on page 390, “trc\_find\_first, trc\_find\_next, and trc\_compare Subroutine” on page 381, “trc\_seek and trc\_tell Subroutine” on page 398, “trc\_strerror Subroutine” on page 399, and “trc\_hookname Subroutine” on page 387.

The perror subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

---

## trc\_read Subroutine

### Purpose

Reads from a trace log object.

### Library

libtrace.a

### Syntax

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

int trc_read (handle, ret)
trc_log_handle_t handle;
trc_read_t *ret;
```

### Description

The **trc\_read** subroutine reads the next sequential data item from the trace log object whose handle is contained in the *handle* parameter. If the **trc\_read** subroutine follows a **trc\_find\_first** or **trc\_find\_next** call, it reads the next sequential data item after the one found. To read the next item matching that criteria, use the **trc\_find\_next** subroutine.

### Parameters

*handle* Contains the handle returned from a successful call to the **trc\_open** subroutine.

*ret* Points to the **trc\_read\_t** structure to contain the returned information. The raw data will be formatted the same way it is formatted today in the **trcpt** internal data buffer. This is described in the **/etc/trcfmt** file for both 32 and 64 bit events. Thus 32-bit trace items will be formatted as 32-bit items regardless of whether they came from a 32 or 64 bit trace. If **TRC\_LOGVERBATIM** was specified, data is returned exactly as traced.

Processed data is the result of trace template processing, see the **/etc/trcfmt** file.

The **trc\_free** subroutine should be used to free data referenced from the **trc\_read\_t** data type. The **trc\_free** subroutine need not be used if the **TRC\_LOGLIVE** flag was specified when the object was opened.

The **/usr/include/sys/libtrace.h** file contains the data definitions for the returned data.

Label	Data Type	Description
<b>trcr_magic</b>	<b>int</b>	Trace read data magic number. This is maintained by the library to identify the library version in use.
<b>trcr_flags</b>	<b>int</b>	Flags that describe the data returned.

The following are definitions for raw data items:

Label	Data Type	Description
<b>trcri_hookid</b>	<b>trc_hookid_t</b>	Trace hook id of the form 0x0hhh, where hhh is the hook id value, (for example) 134.
<b>trcri_subhookid</b>	<b>trc_subhookid_t</b>	Subhook id.
<b>trcri_cpuid</b>	<b>unsigned</b>	The cpu id if known. If the <b>TRCRF_CPUIDOK</b> flag is set, the cpuid value could be determined, otherwise it should be ignored.
<b>trcri_tid</b>	<b>unsigned long long</b>	Thread id.
<b>trcri_timestamp</b>	<b>unsigned long long</b>	Specifies the timestamp in ticks. Use the <b>trc_ticks2nanos</b> function to convert this value to nanoseconds.
<b>trcri_rawofst</b>	<b>unsigned long long</b>	The offset to the start of this trace item in the trace log file.
<b>trcri_rawlen</b>	<b>int</b>	The length of the raw data as traced. This is not necessarily the amount of space used for the data in the log file.

The following are definitions for processed data items:

Label	Data Type	Description
<b>trchi_hookid</b>	<b>trc_hookid_t</b>	The trace hook id of the form 0x0hhh, where hhh is the hook id value, (for example) 134.
<b>trchi_subhookid</b>	<b>trc_subhookid_t</b>	Subhook id.
<b>trchi_elapsed_nseconds</b>	<b>unsigned long long</b>	The elapsed time from the start of the trace in nanoseconds.
<b>trchi_tid</b>	<b>unsigned long long</b>	Thread id.
<b>trchi_pid</b>	<b>unsigned long long</b>	Process id.
<b>trchi_svc</b>	<b>unsigned long long</b>	System call address.
<b>trchi_rawofst</b>	<b>unsigned long long</b>	Offset of the trace event in the log file.
<b>trchi_trcontime</b>	<b>time64_t</b>	The time of the last <b>TRCON</b> , or this <b>TRCON</b> .
<b>trchi_trcofftime</b>	<b>time64_t</b>	The time of the last <b>TRCOFF</b> , or this <b>TRCOFF</b> .
<b>trchi_cpuid</b>	<b>int</b>	cpu id.
<b>trchi_rcpu</b>	<b>int</b>	cpus remaining in this trace.
<b>trchi_pri</b>	<b>int</b>	Process priority.
<b>trchi_intr_depth</b>	<b>int</b>	Interrupt depth.

Label	Data Type	Description
<b>trchi_indent</b>	int	The indentation level used by <b>trcrpt</b> . The values are -1 - <b>\$NOPRINT</b> , 0 - no indentation, 1 - application level, 2 - SVC level, 3 - kernel level. Items greater than zero specify the number of tabs, minus 1, that precede each line of the ascii data, see the <b>trchi_ascii</b> field. Each tab represents 8 blanks, so <b>trchi_indent</b> = 2 implies 2 - 1, or 1 tab before each line of data, or 8 blanks.
<b>trchi_svcname</b>	char *	Current svc name.
<b>trchi_procname</b>	char *	Current process name.
<b>trchi_filename</b>	char *	Current file name.
<b>trchi_ascii</b>	char *	This is the data produced by the trace template for this hook. Each line of data is indented with blanks, according to the <b>trchi_indent</b> value, and the text offset and the subsequent line offset, see the <b>trc_libcntl</b> subroutine.

The **trcr\_flags** field contains bit flags describing characteristics of the returned data. The values are:

<b>TRCRF_RAW</b>	Raw data was read, (for example) the log object was opened with the <b>TRC_LOGRAW</b> open type. Use the raw data items in the return data, (for example) those beginning with <b>trcri_</b> .
<b>TRCRF_PROC</b>	Processed data was read, (for example) the log object was opened with the <b>TRC_LOGPROC</b> open type. Use the processed data items in the return data, (for example) those beginning with <b>trchi_</b> .
<b>TRCRF_64BIT</b>	The data is from a 64-bit environment. Note that the trace itself may be from a 32 or 64 bit kernel.
<b>TRCRF_TIMESTAMPED</b>	The entry was timestamped when traced.
<b>TRCRF_CPUIDOK</b>	The cpu id is known. This is always set for a processed entry, and set for a raw entry if the cpuid was contained in each trace hook (see the <b>-p trace</b> command option), or the trace is a multi-cpu trace (see the <b>-C trace</b> option). For a processed trace, the cpu id may not be accurate if the appropriate hooks, 106 and 10C, weren't traced.
<b>TRCRF_GENERIC</b>	This is a generic trace entry, one traced with the <b>TRCGEN</b> or <b>TRCGENT</b> macros. This is set for a raw trace only.
<b>TRCRF_64BITTRACE</b>	This is a 64-bit trace, (for example) it was taken with a 64-bit kernel.
<b>TRCRF_LIVEDATA</b>	The data is live, don't free it. The data will be changed when another read operation is done.
<b>TRCRF_NOPRINT</b>	The associated trace template specified <b>\$NOPRINT</b> or <b>\$SKIP</b> , (for example) no data should be printed.

## Return Values

Upon successful completion, the **trc\_read** subroutine returns a 0 and puts the data into the *ret* area.

## Error Codes

Upon error, the **trc\_read** subroutine sets the **errno** global variable to a value from **errno.h**, and returns either a value from the **errno.h** file or an error defined in the **libtrace.h** file.

**EINVAL**                   The handle is not valid.  
**TRCE\_BADFORMAT**        The trace data is improperly formatted, and the **errno** global variable is set to **EINVAL**.

## Related Information

“trc\_open Subroutine” on page 392, “trc\_close Subroutine” on page 380, “trc\_loginfo Subroutine” on page 390, “trc\_find\_first, trc\_find\_next, and trc\_compare Subroutine” on page 381, “trc\_libcntl Subroutine” on page 388, “trc\_strerror Subroutine” on page 399, “trc\_perror Subroutine” on page 394, “trcstart Subroutine” on page 405, “trcon Subroutine” on page 404, “trcoff Subroutine” on page 404 and “trcstop Subroutine” on page 406.

The trace daemon in *AIX 5L Version 5.2 Commands Reference, Volume 5*.

The trcrpt, trcstop, and trcupdate commands in *AIX 5L Version 5.2 Commands Reference, Volume 5*.

---

## trc\_seek and trc\_tell Subroutine

### Purpose

Seeks into a trace object and returns the current position that will be used with a future seek.

### Library

**libtrace.a**

### Syntax

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

int trc_seek (handle, log_positionp, r)
trc_loghandle_t handle;
trc_logpos_t log_positionp;
trc_read_t *r;

int trc_tell (handle, log_positionp)
trc_loghandle_t handle;
trc_logpos_t *log_positionp;
```

### Description

The **trc\_seek** subroutine seeks into the log object identified by the *handle* parameter. The *log\_positionp* parameter must have been obtained from a previous call to the **trc\_tell** subroutine. If the **trc\_read\_t** pointer, *r*, is not NULL, the **trc\_seek** subroutine returns the trace data at the seek point.

The **trc\_tell** subroutine creates a **trc\_logpos\_t** object using the current log position and state.

The **trc\_free** subroutine should be used to free a **trc\_logpos\_t** object that's no longer needed. However, **trc\_free** is not necessary if the **trc\_logpos\_t** object is passed to another **trc\_tell**.

## Parameters

<i>handle</i>	Contains the handle returned from a successful call to the <b>trc_open</b> subroutine.
<i>log_positionp</i>	A <b>trc_logpos_t</b> returned by a previous call to the <b>trc_tell</b> subroutine.
<i>r</i>	If not NULL, points to a <b>trc_read_t data</b> item where the data at the new position is returned.

## Return Values

Upon successful return, the **trc\_seek** and **trc\_tell** subroutines return 0.

## Error Codes

If unsuccessful, the **trc\_seek** subroutine returns an i/o error, or **EINVAL** if either the *handle* or *log\_positionp* parameter is in error.

Upon error, the **trc\_tell** subroutine returns **EINVAL** if the handle is invalid, or **ENOMEM** if storage can't be obtained for the **trc\_logpos\_t** object.

## Related Information

“trc\_open Subroutine” on page 392, “trc\_read Subroutine” on page 395, “trc\_loginfo Subroutine” on page 390, “trc\_find\_first, trc\_find\_next, and trc\_compare Subroutine” on page 381, “trc\_libcntl Subroutine” on page 388, “trc\_strerror Subroutine”, “trc\_perror Subroutine” on page 394, and “trc\_hookname Subroutine” on page 387.

The perror subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

---

## trc\_strerror Subroutine

### Purpose

Returns the error message, or next error message, associated with a trace log object or **trc\_loginfo** object.

### Library

libtrace.a

### Syntax

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

```
char *trc_strerror (handle, rv)  
void *handle;  
int rv;
```

### Description

The **trc\_strerror** subroutine is similar to the **strerror** subroutine. If the error in the *rv* parameter is an error from the **errno.h** file, it simply returns the string from the **strerror** subroutine. If the *rv* parameter is a **libtrace** error such as **TRCE\_EOF**, it returns the string associated with this error. It is possible for multiple **libtrace** errors to be present. The **trc\_strerror** subroutine returns the next error in this case. When no more errors are present, the **trc\_strerror** subroutine returns NULL.

Like the **strerror** subroutine, the **trc\_strerror** subroutine must not be used in a threaded environment.

## Parameters

- handle* Contains the handle returned from the **trc\_open** subroutine, the pointer to a **trc\_loginfo\_t** object, or NULL. If a handle returned by the **trc\_open** subroutine is passed, the **trc\_open** subroutine need not have been successful, but the **TRC\_RETAIN\_HANDLE** open option must have been used.
- rv* Contains the return value from a call to the **libtrace** subroutine.

## Return Values

The **trc\_strerror** subroutine returns a pointer to the associated error message. It returns NULL if no more errors are present.

## Examples

1. To retrieve all error messages from a call to the **trc\_open** subroutine, call the **trc\_strerror** subroutine as follows:

```
{
    trc_loghandle_t h;
    int rv;
    char *fn, *tfn, *s;

    ...

    rv = trc_open(fn,tfn, TRC_LOGREAD|TRC_LOGPROC|TRC_RETAIN_HANDLE, &h);
    while (rv && s=trc_strerror(h, rv)) {
        fprintf(stderr, "%s\n", s);
    }
}
```

2. To accomplish the same thing as the previous example with a single call, do the following:

```
{
    trc_loghandle_t h;
    int rv;
    char *fn, *tfn;

    ...

    rv = trc_open(fn,tfn, TRC_LOGREAD|TRC_LOGPROC|TRC_RETAIN_HANDLE, &h);
    if (rv) trc_perror(h, rv, "");
}
```

## Related Information

“trc\_open Subroutine” on page 392, “trc\_read Subroutine” on page 395, “trc\_loginfo Subroutine” on page 390, “trc\_find\_first, trc\_find\_next, and trc\_compare Subroutine” on page 381, “trc\_seek and trc\_tell Subroutine” on page 398, “trc\_perror Subroutine” on page 394, “trc\_hookname Subroutine” on page 387, and “strerror Subroutine” on page 285.

---

## trcgen or trcgent Subroutine

### Purpose

Records a trace event for a generic trace channel.

### Library

Runtime Services Library (**librts.a**)

### Syntax

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

```

void trcgen(Channel, HkWord, DataWord, Length, Buffer)
unsigned int Channel, HkWord, DataWord, Length;
char * Buffer;

void trcgent(Channel, HkWord, DataWord, Length, Buffer)
unsigned int Channel, HkWord, DataWord, Length;
char *Buffer;

```

## Description

The **trcgen** subroutine records a trace event for a generic trace entry consisting of a hook word, a data word, and a variable number of bytes of trace data. The **trcgent** subroutine records a trace event for a generic trace entry consisting of a hook word, a data word, a variable number of bytes of trace data, and a time stamp.

The **trcgen** subroutine and **trcgent** subroutine are located in pinned kernel memory.

## Parameters

*Buffer* Specifies a pointer to a buffer of trace data. The maximum size of the trace data is 4096 bytes.

*Channel* Specifies a channel number for the trace session, obtained from the **trcstart** subroutine.

*DataWord* Specifies a word of user-defined data.

*HkWord* Specifies an integer consisting of two bytes of user-defined data (*HkData*), a hook ID (*HkID*), and a hook type (*Hk\_Type*).

*HkData* Specifies two bytes of user-defined data.

*HkID* Specifies a hook identifier. For user programs, the hook ID value ranges from 010 to 0FF.

*Hk\_Type* Specifies a 4-bit value that identifies the amount of trace data to be recorded:

Value	Records
1	Hook word
9	Hook word and a time stamp
2	Hook word and one data word
A	Hook word, one data word, and a time stamp
6	Hook word and up to five data words
E	Hook word, up to five data words, and a time stamp.

*Length* Specifies the length in bytes of the *Buffer* parameter.

## Related Information

The **trchook** (“trchook, utrchook, trchook64, and utrhook64 Subroutine” on page 402) subroutine, **trcoff** (“trcoff Subroutine” on page 404) subroutine, **trcon** (“trcon Subroutine” on page 404) subroutine, **trcstart** (“trcstart Subroutine” on page 405) subroutine, **trcstop** (“trcstop Subroutine” on page 406) subroutine.

The **trace** daemon.

The **trcgenk** kernel service, **trcgenkt** kernel service.

---

## trchook, utrchook, trchook64, and utrhook64 Subroutine

### Purpose

Records a trace event.

### Library

Runtime Services Library (**librts.a**)

### Syntax

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

```
void trchook( HkWord, d1, d2, d3, d4, d5)
```

```
unsigned int HkWord, d1, d2, d3, d4, d5;
```

```
void utrchook(HkWord, d1, d2, d3, d4, d5)
```

```
unsigned int HkWord, d1, d2, d3, d4, d5;
```

```
void trchook64 (HkWord, d1, d2, d3, d4, d5)
```

```
unsigned long HkWord, d1, d2, d3, d4, d5;
```

```
void utrchook64 (HkWord, d1, d2, d3, d4, d5)
```

```
unsigned long HkWord, d1, d2, d3, d4, d5;
```

### Description

The **trchook** subroutine records a trace event if a trace session is active. Input parameters include a hook word (*HkWord*) and from 0 to 5 words of data. The **trchook** and **trchook64** subroutines are intended for use by the kernel and extensions.

The **utrchook** and **utrchook64** subroutines are intended for programs running at user (application) level.

The **trchook** and **utrchook** subroutines are for use in a 32-bit environment, while the **trchook64** and **utrchook64** subroutines are intended for use in a 64-bit environment. Note that if running a 64-bit application on a 32-bit kernel, the application should use **utrchook64**(the subroutine for its 64-bit environment).

It is strongly recommended that the C macros **TRCHKLn** and **TRCHKLnT** (where **n** is from 0 to 5) be used if possible, instead of calling these subroutines directly.

### Parameters

*d1, d2, d3, d4, d5*

Up to 5 words of data from the calling program.

*HkWord*

The *HkWord* parameter has a different format based upon the environment. For the **trchook** and **utrchook** subroutines, it is an unsigned long consisting of a hook ID (*HkID*), a hook type (*Hk\_Type*), and two bytes of data from the calling program (*HkData*).

*HkID* A hook ID is a 12-bit value. For user programs, the hook ID may be a value from 0x010 to 0x0FF. Hook identifiers are defined in the **/usr/include/sys/trchkid.h** file.

*Hk\_Type*

A 4-bit value that identifies the amount of trace data to be recorded:

Value	Records
1	Hook word
9	Hook word and a time stamp
2	Hook word and one data word
A	Hook word, one data word, and a time stamp
6	Hook word and up to five data words
E	Hook word, up to five data words, and a time stamp.

*HkData* Two bytes of data from the calling program.

In a 64-bit environment, when using the **trchook64** or **utrchook64** subroutine, the format is *ffffllllhhhxssss*, where *f* represents flags, *l* is length, *h* is the hook id, and *s* is the subhook.

The hook and subhook ids are the same as for the 32-bit environment (12-bit hook id and a 16-bit subhook id). Note that the 4 bits between the hook id and subhook are unused.

The flags (the first 16 bits of the 64-bit hookword) are specified as follows:

- 8000** The hook should be timestamped.
- 4000** A generic trace entry, should not use the **trchook64** or **utrchook64** subroutine. For more information see “trcgen or trcgent Subroutine” on page 400.
- 2000** The hook contains 32-bit data. Used by aix trace only.
- 1000** Automatically include the cpuid when tracing the data.

The length (*l*) is the second 16 bits of the hookword. It is the length of the data. The length is 0 if no data other than the hookword is traced (**TRCHKL0**), 8 if one parameter, 8 bytes, is traced (**TRCHKL1**), 16 for 2 parameters, 24 for 3 parameters, 32 for 4 parameters, and 40 for 5 parameters (**TRCHKL5**).

## Related Information

The **trcgen** (“trcgen or trcgent Subroutine” on page 400) subroutine, **trcgent** (“trcgen or trcgent Subroutine” on page 400) subroutine, **trcoff** (“trcoff Subroutine” on page 404) subroutine, **trcon** (“trcon Subroutine” on page 404) subroutine, **trcstart** (“trcstart Subroutine” on page 405) subroutine, **trcstop** (“trcstop Subroutine” on page 406) subroutine.

The **trace** daemon.

The **trcgenk** kernel service, **trcgenkt** kernel service.

---

## trcoff Subroutine

### Purpose

Halts the collection of trace data from within a process.

### Library

Runtime Services Library (**librts.a**)

### Syntax

```
int trcoff( Channel)
int Channel;
```

### Description

The **trcoff** subroutine stops trace data collection for a trace channel. The trace session must have already been started using the **trace** command or the **trcstart** subroutine.

### Parameters

*Channel*      Channel number for the trace session.

### Return Values

If the **trcoff** subroutine was successful, zero is returned and trace data collection stops. If unsuccessful, a negative one is returned.

### Related Information

The **trcgen** (“trcgen or trcgent Subroutine” on page 400) subroutine, **trchook** (“trchook, utrchook, trchook64, and utrhook64 Subroutine” on page 402) subroutine, **trcon** (“trcon Subroutine”) subroutine, **trcstart** (“trcstart Subroutine” on page 405) subroutine, **trcstop** (“trcstop Subroutine” on page 406) subroutine.

The **trace** daemon.

**trcgenk** kernel service, **trcgenkt** kernel service.

---

## trcon Subroutine

### Purpose

Starts the collection of trace data.

### Library

Runtime Services Library (**librts.a**)

### Syntax

```
int trcon( Channel)
int Channel;
```

### Description

The **trcon** subroutine starts trace data collection for a trace channel. The trace session must have already been started using the **trace** command or the **trcstart** (“trcstart Subroutine” on page 405) subroutine.

## Parameters

*Channel* Specifies one of eight trace channels. Channel number 0 always refers to the Event/Performance trace. Channel numbers 1 through 7 specify generic trace channels.

## Return Values

If the **trcon** subroutine was successful, zero is returned and trace data collection starts. If unsuccessful, a negative one is returned.

## Related Information

The **trcgen** (“trcgen or trcgent Subroutine” on page 400) subroutine, **trchook** (“trchook, utrchook, trchook64, and utrhook64 Subroutine” on page 402) subroutine, **trcoff** (“trcoff Subroutine” on page 404) subroutine, **trcstart** (“trcstart Subroutine”) subroutine, **trcstop** (“trcstop Subroutine” on page 406) subroutine.

The **trace** daemon.

The **trcgenk** kernel service, **trcgenkt** kernel service.

---

## trcstart Subroutine

### Purpose

Starts a trace session.

### Library

Runtime Services Library (**librts.a**)

### Syntax

```
int trcstart( Argument )
char *Argument;
```

### Description

The **trcstart** subroutine starts a trace session. The *Argument* parameter points to a character string containing the flags invoked with the **trace** daemon. To specify that a generic trace session is to be started, include the **-g** flag.

### Parameters

*Argument* Character pointer to a string holding valid arguments from the **trace** daemon.

### Return Values

If the **trace** daemon is started successfully, the channel number is returned. Channel number 0 is returned if a generic trace was not requested. If the **trace** daemon is not started successfully, a value of -1 is returned.

### Files

**/dev/trace** Trace special file.

## Related Information

The **trcon** (“trcon Subroutine” on page 404) subroutine.

The **trace** daemon.

---

## trcstop Subroutine

### Purpose

Stops a trace session.

### Library

Runtime Services Library (**librts.a**)

### Syntax

```
int trcstop( Channel)
int Channel;
```

### Description

The **trcstop** subroutine stops a trace session for a particular trace channel.

### Parameters

*Channel* Specifies one of eight trace channels. Channel number 0 always refers to the Event/Performance trace. Channel numbers 1 through 7 specify generic trace channels.

### Return Values

0 The trace session was stopped successfully.  
-1 The trace session did not stop.

## Related Information

The **trcgen** (“trcgen or trcgent Subroutine” on page 400) subroutine, **trchhook** (“trchhook, utrchhook, trchhook64, and utrhook64 Subroutine” on page 402) subroutine, **trcoff** (“trcoff Subroutine” on page 404) subroutine, **trcon** (“trcon Subroutine” on page 404) subroutine, **trcstart** (“trcstart Subroutine” on page 405) subroutine.

The **trace** daemon.

The **trcgenk** kernel service, **trcgenkt** kernel service.

---

## trunc, truncf, or trunci Subroutine

### Purpose

Rounds to truncated integer value.

### Syntax

```
#include <math.h>
double trunc (x)
```

**double** *x*;

**float truncf** (*x*)

**float** *x*;

**long double trunc1** (*x*)

**long double** *x*;

## Description

The **trunc**, **truncf**, and **trunc1** subroutines round the *x* parameter to the integer value, in floating format, nearest to but no larger in magnitude than the *x* parameter.

## Parameters

*x* Specifies the value to be rounded.

## Return Values

Upon successful completion, the **trunc**, **truncf**, and **trunc1** subroutines return the truncated integer value.

If *x* is NaN, a NaN is returned.

If *x* is  $\pm 0$  or  $\pm \text{Inf}$ , *x* is returned.

## Related Information

**math.h** in *AIX 5L Version 5.2 Files Reference*.

---

## truncate, truncate64, ftruncate, or ftruncate64 Subroutine

### Purpose

Changes the length of regular files.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int truncate ( Path, Length)
```

```
const char *Path;
```

```
off_t Length;
```

```
int ftruncate ( FileDescriptor, Length)
```

```
int FileDescriptor;
```

```
off_t Length;
```

**Note:** The **truncate64** and **ftruncate64** subroutines apply to AIX 4.2 and later releases.

```
int truncate64 ( Path, Length)
```

```
const char *Path;
```

```
off64_t Length;
```

```
int ftruncate64 ( FileDescriptor, Length)
int FileDescriptor;
off64_t Length;
```

## Description

**Note:** The **truncate64** and **ftruncate64** subroutines apply to AIX 4.2 and later releases.

The **truncate** and **ftruncate** subroutines change the length of regular files.

The *Path* parameter must point to a regular file for which the calling process has write permission. The *Length* parameter specifies the desired length of the new file in bytes.

The *Length* parameter measures the specified file in bytes from the beginning of the file. If the new length is less than the previous length, all data between the new length and the previous end of file is removed. If the new length in the specified file is greater than the previous length, data between the old and new lengths is read as zeros. Full blocks are returned to the file system so that they can be used again, and the file size is changed to the value of the *Length* parameter.

If the file designated in the *Path* parameter names a symbolic link, the link will be traversed and path-name resolution will continue.

These subroutines do not modify the seek pointer of the file.

These subroutines cannot be applied to a file that a process has open with the **O\_DEFER** flag.

Successful completion of the **truncate** or **ftruncate** subroutine updates the *st\_ctime* and *st\_mtime* fields of the file. Successful completion also clears the SetUserID bit (**S\_ISUID**) of the file if any of the following are true:

- The calling process does not have root user authority.
- The effective user ID of the calling process does not match the user ID of the file.
- The file is executable by the group (**S\_IXGRP**) or others (**S\_IXOTH**).

These subroutines also clear the SetGroupID bit (**S\_ISGID**) if:

- The file does not match the effective group ID or one of the supplementary group IDs of the process
- OR
- The file is executable by the owner (**S\_IXUSR**) or others (**S\_IXOTH**).

**Note:** Clearing of the SetUserID and SetGroupID bits can occur even if the subroutine fails because the data in the file was modified before the error was detected.

**truncate** and **ftruncate** can be used to specify any size up to **OFF\_MAX**. **truncate64** and **ftruncate64** can be used to specify any length up to the maximum file size for the file.

In the large file enabled programming environment, **truncate** is redefined to be **truncate64** and **ftruncate** is redefined to be **ftruncate64**.

## Parameters

<i>Path</i>	Specifies the name of a file that is opened, truncated, and then closed.
<i>FileDescriptor</i>	Specifies the descriptor of a file that must be open for writing.
<i>Length</i>	Specifies the new length of the truncated file in bytes.

## Return Values

Upon successful completion, a value of 0 is returned. If the **truncate** or **ftruncate** subroutine is unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the nature of the error.

## Error Codes

The **truncate** and **ftruncate** subroutines fail if the following is true:

**EROFS** An attempt was made to truncate a file that resides on a read-only file system.

**Note:** In addition, the **truncate** subroutine can return the same errors as the **open** subroutine if there is a problem opening the file.

The **truncate** and **ftruncate** subroutines fail if one of the following is true:

<b>EAGAIN</b>	The truncation operation fails due to an enforced write lock on a portion of the file being truncated. Because the target file was opened with the <b>O_NONBLOCK</b> or <b>O_NDELAY</b> flags set, the subroutine fails immediately rather than wait for a release.
<b>EDQUOT</b>	New disk blocks cannot be allocated for the truncated file. The quota of the user's or group's allotted disk blocks has been exhausted on the target file system.
<b>EFBIG</b>	An attempt was made to write a file that exceeds the process' file size limit or the maximum file size. If the user has set the environment variable <b>XPG_SUS_ENV=ON</b> prior to execution of the process, then the <b>SIGXFSZ</b> signal is posted to the process when exceeding the process' file size limit.
<b>EFBIG</b>	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>files</i> .
<b>EINVAL</b>	The file is not a regular file.
<b>EINVAL</b>	The <i>Length</i> parameter was less than zero.
<b>EISDIR</b>	The named file is a directory.
<b>EINTR</b>	A signal was caught during execution.
<b>EIO</b>	An I/O error occurred while reading from or writing to the file system.
<b>EMFILE</b>	The file is open with <b>O_DEFER</b> by one or more processes.
<b>ENOSPC</b>	New disk blocks cannot be allocated for the truncated file. There is no free space on the file system containing the file.
<b>ETXTBSY</b>	The file is part of a process that is running.
<b>EROFS</b>	The named file resides on a read-only file system.

### Notes:

1. The **truncate** subroutine can also be unsuccessful for other reasons. For a list of additional errors, see "Base Operating System Error Codes For Services That Require Path-Name Resolution" .
2. The **truncate** subroutine can return the same errors as the **open** subroutine if there is a problem opening the file.

The **ftruncate** subroutine fails if the following is true:

<b>EBADF</b>	The <i>FileDescriptor</i> parameter is not a valid file descriptor open for writing.
<b>EINVAL</b>	The <i>FileDescriptor</i> argument references a file that was opened without write permission.

The **truncate** function will fail if:

<b>EACCES</b>	A component of the path prefix denies search permission, or write permission is denied on the file.
<b>EISDIR</b>	The named file is a directory.
<b>ELOOP</b>	Too many symbolic links were encountered in resolving <i>path</i> .

<b>ENAMETOOLONG</b>	The length of the specified pathname exceeds <b>PATH_MAX</b> bytes, or the length of a component of the pathname exceeds <b>NAME_MAX</b> bytes.
<b>ENOENT</b>	A component of <i>path</i> does not name an existing file or <i>path</i> is an empty string.
<b>ENTDIR</b>	A component of the path prefix of <i>path</i> is not a directory.
<b>EROFS</b>	The named file resides on a read-only file system.

The **truncate** function may fail if:

<b>ENAMETOOLONG</b>	Pathname resolution of a symbolic link produced an intermediate result whose length exceeds <b>PATH_MAX</b> .
---------------------	---

## Related Information

The **fclear** subroutine, **openx**, **open**, or **creat** subroutine.

Appendix A, “Base Operating System Error Codes for Services That Require Path-Name Resolution”, on page 721.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tsearch, tdelete, tfind or twalk Subroutine

### Purpose

Manages binary search trees.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <search.h>
```

```
void *tsearch ( Key, RootPointer, ComparisonPointer)
const void *Key;
void **RootPointer;
int (*ComparisonPointer) (const void *Element1, const void *Element2);
```

```
void *tdelete (Key, RootPointer, ComparisonPointer)
const void *Key;
void **RootPointer;
int (*ComparisonPointer) (const void *Element1, const void *Element2);
```

```
void *tfind (Key, RootPointer, ComparisonPointer)
const void *Key;
void *const *RootPointer;
int (*ComparisonPointer) (const void *Element1, const void *Element2);
```

```
void twalk ( Root, Action)
const void *Root;
void (*Action) (const void *Node, VISIT Type, int Level);
```

## Description

The **tsearch**, **tdelete**, **tfind** and **twalk** subroutines manipulate binary search trees. Comparisons are made with the user-supplied routine specified by the *ComparisonPointer* parameter. This routine is called with two parameters, the pointers to the elements being compared.

The **tsearch** subroutine performs a binary tree search, returning a pointer into a tree indicating where the data specified by the *Key* parameter can be found. If the data specified by the *Key* parameter is not found, the data is added to the tree in the correct place. If there is not enough space available to create a new node, a null pointer is returned. Only pointers are copied, so the calling routine must store the data. The *RootPointer* parameter points to a variable that points to the root of the tree. If the *RootPointer* parameter is the null value, the variable is set to point to the root of a new tree. If the *RootPointer* parameter is the null value on entry, then a null pointer is returned.

The **tdelete** subroutine deletes the data specified by the *Key* parameter. The *RootPointer* and *ComparisonPointer* parameters perform the same function as they do for the **tsearch** subroutine. The variable pointed to by the *RootPointer* parameter is changed if the deleted node is the root of the binary tree. The **tdelete** subroutine returns a pointer to the parent node of the deleted node. If the data is not found, a null pointer is returned. If the *RootPointer* parameter is null on entry, then a null pointer is returned.

The **tfind** subroutine searches the binary search tree. Like the **tsearch** subroutine, the **tfind** subroutine searches for a node in the tree, returning a pointer to it if found. However, if it is not found, the **tfind** subroutine will return a null pointer. The parameters for the **tfind** subroutine are the same as for the **tsearch** subroutine.

The **twalk** subroutine steps through the binary search tree whose root is pointed to by the *RootPointer* parameter. (Any node in a tree can be used as the root to step through the tree below that node.) The *Action* parameter is the name of a routine to be invoked at each node. The routine specified by the *Action* parameter is called with three parameters. The first parameter is the address of the node currently being pointed to. The second parameter is a value from an enumeration data type:

```
typedef enum [preorder, postorder, endorder, leaf] VISIT;
```

(This data type is defined in the **search.h** file.) The actual value of the second parameter depends 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*. A leaf is a node that is not the parent of another node. The third parameter is the level of the node in the tree, with the root node being level zero.

Although declared as type pointer-to-void, 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. Although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

## Parameters

<i>Key</i>	Points to the data to be located.
<i>ComparisonPointer</i>	Points to the comparison function, which is called with two parameters that point to the elements being compared.
<i>RootPointer</i>	Points to a variable that in turn points to the root of the tree.
<i>Action</i>	Names a routine to be invoked at each node.
<i>Root</i>	Points to the roots of a binary search node.

## Return Values

The comparison function compares its parameters and returns a value as follows:

- If the first parameter is less than the second parameter, the *ComparisonPointer* parameter returns a value less than 0.

- If the first parameter is equal to the second parameter, the *ComparisonPointer* parameter returns a value of 0.
- If the first parameter is greater than the second parameter, the *ComparisonPointer* parameter returns a value greater than 0.

The comparison function need not compare every byte, so arbitrary data can be contained in the elements in addition to the values being compared.

If the node is found, the **tsearch** and **tfind** subroutines return a pointer to it. If the node is not found, the **tsearch** subroutine returns a pointer to the inserted item and the **tfind** subroutine returns a null pointer. If there is not enough space to create a new node, the **tsearch** subroutine returns a null pointer.

If the *RootPointer* parameter is a null pointer on entry, a null pointer is returned by the **tsearch** and **tdelete** subroutines.

The **tdelete** subroutine returns a pointer to the parent of the deleted node. If the node is not found, a null pointer is returned.

## Related Information

The **bsearch** subroutine, **hsearch** subroutine, **lsearch** subroutine.

Searching and Sorting Example Program, Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## ttylock, ttywait, ttyunlock, or ttylocked Subroutine

### Purpose

Controls tty locking functions.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int ttylock ( DeviceName)
char *DeviceName;

int ttywait (DeviceName)
char *DeviceName;

int ttyunlock (DeviceName)
char *DeviceName;

int ttylocked (DeviceName)
char *DeviceName;
```

### Description

The **ttylock** subroutine creates the **LCK..DeviceName** file in the **/etc/locks** directory and writes the process ID of the calling process in that file. If **LCK..DeviceName** exists and the process whose ID is contained in this file is active, the **ttylock** subroutine returns an error.

There are programs like **uucp** and **connect** that create tty locks in the **/etc/locks** directory. The convention followed by these programs is to call the **ttylock** subroutine with an argument of *DeviceName* for locking the **/dev/DeviceName** file. This convention must be followed by all callers of the **ttylock** subroutine to make the locking mechanism work.

The **ttymwait** subroutine blocks the calling process until the lock file associated with *DeviceName*, the */etc/locks/LCK..DeviceName* file, is removed.

The **ttymunlock** subroutine removes the lock file, */etc/locks/LCK..DeviceName*, if it is held by the current process.

The **ttymlocked** subroutine checks to see if the lock file, */etc/locks/LCK..DeviceName*, exists and the process that created the lock file is still active. If the process is no longer active, the lock file is removed.

## Parameters

*DeviceName*            Specifies the name of the device.

## Return Values

Upon successful completion, the **ttymlock** subroutine returns a value of 0. Otherwise, a value of -1 is returned.

The **ttymlocked** subroutine returns a value of 0 if no process has a lock on device. Otherwise, a value of -1 is returned.

## Examples

1. To create a lock for */dev/tty0*, use the following statement:

```
rc = ttymlock("tty0");
```

2. To lock */dev/tty0* device and wait for lock to be cleared if it exists, use the following statements:

```
if (ttymlock("tty0"))
    ttymwait("tty0");
rc = ttymlock("tty0");
```

3. To remove the lock file for device */dev/tty0* created by a previous call to the **ttymlock** subroutine, use the following statement:

```
ttymunlock("tty0");
```

4. To check for a lock on */dev/tty0*, use the following statement:

```
rc = ttymlocked("tty0");
```

## Related Information

The */etc/locks* directory.

The Input and Output Handling Programmer's Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## ttymname or isatty Subroutine

### Purpose

Gets the name of a terminal or determines if the device is a terminal.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
char *ttyname( FileDescriptor)
int FileDescriptor;
int isatty(FileDescriptor)
int FileDescriptor;
```

## Description

**Attention:** Do not use the **ttyname** subroutine in a multithreaded environment.

The **ttyname** subroutine gets the path name of a terminal.

The **isatty** subroutine determines if the file descriptor specified by the *FileDescriptor* parameter is associated with a terminal.

The **isatty** subroutine does not necessarily indicate that a person is available for interaction, since nonterminal devices may be connected to the communications line.

## Parameters

*FileDescriptor*                      Specifies an open file descriptor.

## Return Values

The **ttyname** subroutine returns a pointer to a string containing the null-terminated path name of the terminal device associated with the file descriptor specified by the *FileDescriptor* parameter. A null pointer is returned and the **errno** global variable is set to indicate the error if the file descriptor does not describe a terminal device in the **/dev** directory.

The return value of the **ttyname** subroutine may point to static data whose content is overwritten by each call.

If the specified file descriptor is associated with a terminal, the **isatty** subroutine returns a value of 1. If the file descriptor is not associated with a terminal, a value of 0 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **ttyname** and **isatty** subroutines are unsuccessful if one of the following is true:

**EBADF**        The *FileDescriptor* parameter does not specify a valid file descriptor.  
**ENOTTY**        The *FileDescriptor* parameter does not specify a terminal device.

## Files

**/dev/\***        Terminal device special files.

## Related Information

The **ttyslot** (“ttyslot Subroutine” on page 415) subroutine.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## ttyslot Subroutine

### Purpose

Finds the slot in the **utmp** file for the current user.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
int ttyslot (void)
```

### Description

The **ttyslot** subroutine returns the index of the current user's entry in the **/etc/utmp** file. The **ttyslot** subroutine scans the **/etc/utmp** file for the name of the terminal associated with the standard input, the standard output, or the error output file descriptors (0, 1, or 2).

The **ttyslot** subroutine returns -1 if an error is encountered while searching for the terminal name, or if none of the first three file descriptors (0, 1, and 2) is associated with a terminal device.

### Files

<b>/etc/inittab</b>	The path to the <b>inittab</b> file, which controls the initialization process.
<b>/etc/utmp</b>	The path to the <b>utmp</b> file, which contains a record of users logged in to the system.

### Related Information

The **getutent** subroutine, **ttyname** or **isatty** (“**ttyname** or **isatty** Subroutine” on page 413) subroutine.

The Input and Output Handling Programmer's Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## ulimit Subroutine

### Purpose

Sets and gets user limits.

### Library

Standard C Library (**libc.a**)

### Syntax

The syntax for the **ulimit** subroutine when the *Command* parameter specifies a value of **GET\_FSIZE** or **SET\_FSIZE** is:

```
#include <ulimit.h>
```

```
long int ulimit ( Command, NewLimit )
int Command;
off_t NewLimit;
```

The syntax for the **ulimit** subroutine when the *Command* parameter specifies a value of **GET\_DATALIM**, **SET\_DATALIM**, **GET\_STACKLIM**, **SET\_STACKLIM**, **GET\_REALDIR**, or **SET\_REALDIR** is:

```
#include <ulimit.h>
long int ulimit (Command, NewLimit)
int Command;
unsigned long NewLimit;
```

## Description

The **ulimit** subroutine controls process limits.

Even with remote files, the **ulimit** subroutine values of the process on the client node are used.

**Note:** Raising the data ulimit does not necessarily raise the program break value. If the proper memory segments are not initialized at program load time, raising your memory limit will not allow access to this memory. Also, without these memory segments initialized, the value returned after such a change may not be the proper break value. If your data limit is RLIM\_INFINITY, this value will never advance past the segment size, even if that data is available. Use the **-bmaxdata** flag of the **ld** command to set up these segments at load time.

## Parameters

*Command* Specifies the form of control. The following *Command* parameter values require that the *NewLimit* parameter be declared as an **off\_t** structure:

### GET\_FSIZE (1)

Returns the process file size limit. The limit is in units of **UBSIZE** blocks (see the **sys/param.h** file) and is inherited by child processes. Files of any size can be read. The process file size limit is returned in the **off\_t** structure specified by the *NewLimit* parameter.

### SET\_FSIZE (2)

Sets the process file size limit to the value in the **off\_t** structure specified by the *NewLimit* parameter. Any process can decrease this limit, but only a process with root user authority can increase the limit. The new file size limit is returned.

The following *Command* parameter values require that the *NewLimit* parameter be declared as an integer:

### GET\_DATALIM (3)

Returns the maximum possible break value (as described in the **brk** or **sbrk** subroutine).

### SET\_DATALIM (1004)

Sets the maximum possible break value (described in the **brk** and **sbrk** subroutines). Returns the new maximum break value, which is the *NewLimit* parameter rounded up to the nearest page boundary.

### GET\_STACKLIM (1005)

Returns the lowest valid stack address.  
**Note:** Stacks grow from high addresses to low addresses.

### SET\_STACKLIM (1006)

Sets the lowest valid stack address. Returns the new minimum valid stack address, which is the *NewLimit* parameter rounded down to the nearest page boundary.

### GET\_REALDIR (1007)

Returns the current value of the **real directory read** flag. If this flag is a value of 0, a **read** system call (or **readx** with *Extension* parameter value of 0) against a directory returns fixed-format entries compatible with the System V UNIX operating system. Otherwise, a **read** system call (or **readx** with *Extension* parameter value of 0) against a directory returns the underlying physical format.

### SET\_REALDIR (1008)

Sets the value of the **real directory read** flag. If the *NewLimit* parameter is a value of 0, this flag is cleared; otherwise, it is set. The old value of the **real directory read** flag is returned.

*NewLimit* Specifies the new limit. The value and data type or structure of the *NewLimit* parameter depends on the *Command* parameter value that is used.

## Examples

To increase the size of the stack by 4096 bytes (use 4096 or PAGESIZE), and set the *rc* to the new lowest valid stack address, enter:

```
rc = ulimit(SET_STACKLIM, ulimit(GET_STACKLIM, 0) - 4096);
```

## Return Values

Upon successful completion, the value of the requested limit is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

All return values are permissible if the **ulimit** subroutine is successful. To check for error situations, an application should set the **errno** global variable to 0 before calling the **ulimit** subroutine. If the **ulimit** subroutine returns a value of -1, the application should check the **errno** global variable to verify that it is nonzero.

## Error Codes

The **ulimit** subroutine is unsuccessful and the limit remains unchanged if one of the following is true:

**EPERM** A process without root user authority attempts to increase the file size limit.  
**EINVAL** The *Command* parameter is a value other than **GET\_FSIZE**, **SET\_FSIZE**, **GET\_DATALIM**, **SET\_DATALIM**, **GET\_STACKLIM**, **SET\_STACKLIM**, **GET\_REALDIR**, or **SET\_REALDIR**.

## Related Information

The **brk** subroutine, **sbrk** subroutine, **getrlimit** or **setrlimit** subroutine, **pathconf** subroutine, **read** (“read, readx, readv, readvx, or pread Subroutine” on page 16) subroutines, **vlimit** subroutine, **write** (“write, writex, writev, writevx or pwrite Subroutines” on page 506) subroutine.

---

## umask Subroutine

### Purpose

Sets and gets the value of the file creation mask.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
mode_t umask ( CreationMask )  
mode_t CreationMask;
```

### Description

The **umask** subroutine sets the file-mode creation mask of the process to the value of the *CreationMask* parameter and returns the previous value of the mask.

Whenever a file is created (by the **open**, **mkdir**, or **mknod** subroutine), all file permission bits set in the file mode creation mask are cleared in the mode of the created file. This clearing allows users to restrict the default access to their files.

The mask is inherited by child processes.

### Parameters

*CreationMask* Specifies the value of the file mode creation mask. The *CreationMask* parameter is constructed by logically ORing file permission bits defined in the **sys/mode.h** file. Nine bits of the *CreationMask* parameter are significant.

### Return Values

If successful, the file permission bits returned by the **umask** subroutine are the previous value of the file-mode creation mask. The *CreationMask* parameter can be set to this value in subsequent calls to the **umask** subroutine, returning the mask to its initial state.

## Related Information

The **chmod** subroutine, **mkdir** subroutine, **mkfifo** subroutine, **mknod** subroutine, **openx**, **open**, or **creat** subroutine, **stat** (“statx, stat, lstat, fstatx, fstat, fullstat, ffullstat, stat64, lstat64, or fstat64 Subroutine” on page 277) subroutine.

The **sh** command, **ksh** command.

The **sys/mode.h** file.

Shells Overview in *AIX 5L Version 5.2 System User's Guide: Operating System and Devices*.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## umount or uvmount Subroutine

### Purpose

Removes a virtual file system from the file tree.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int umount ( Device)
char *Device;
#include <sys/vmount.h>
```

```
int uvmount ( VirtualFileSystemID, Flag)
int VirtualFileSystemID;
int Flag;
```

### Description

The **umount** and **uvmount** subroutines remove a virtual file system (VFS) from the file tree.

The **umount** subroutine unmounts only file systems mounted from a block device (a special file identified by its path to the block device).

In addition to local devices, the **uvmount** subroutine unmounts local or remote directories, identified by the *VirtualFileSystemID* parameter.

Only a calling process with root user authority or in the system group and having write access to the mount point can unmount a device, file and directory mount.

### Parameters

*Device*

The path name of the block device to be unmounted for the **umount** subroutine.

*VirtualFileSystemID*

The unique identifier of the VFS to be unmounted for the **uvmount** subroutine. This value is returned when a VFS is created by the **vmount** subroutine and may subsequently be obtained by the **mntctl** subroutine. The *VirtualFileSystemID* is also reported in the **stat** subroutine `st_vfs` field.

Flag

Specifies special action for the **uvmount** subroutine. Currently only one value is defined:

**UVMNT\_FORCE**

Force the unmount. This flag is ignored for device mounts.

## Return Values

Upon successful completion a value of 0 is returned. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **uvmount** subroutine fails if one of the following is true:

- EPERM** The calling process does not have write permission to the root of the VFS, the mounted object is a device or remote, and the calling process does not have root user authority.
- EINVAL** No VFS with the specified *VirtualFileSystemID* parameter exists.
- EBUSY** A device that is still in use is being unmounted.

The **umount** subroutine fails if one of the following is true:

- EPERM** The calling process does not have root user authority.
- ENOENT** The *Device* parameter does not exist.
- ENOBLOCK** The *Device* parameter is not a block device.
- EINVAL** The *Device* parameter is not mounted.
- EINVAL** The *Device* parameter is not local.
- EBUSY** A process is holding a reference to a file located on the file system.

The **umount** subroutine can be unsuccessful for other reasons. For a list of additional errors, see "Base Operating System Error Codes For Services That Require Path-Name Resolution".

The **umount** subroutine can be unsuccessful for other reasons. For a list of additional errors, see Appendix A, "Base Operating System Error Codes for Services That Require Path-Name Resolution."

## Related Information

The **mount** ("vmount or mount Subroutine" on page 435) subroutine.

The **mount** command, **umount** command.

Mounting Overview in *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## uname or unamex Subroutine

### Purpose

Gets the name of the current operating system.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <sys/utsname.h>
int uname ( Name)
struct utsname *Name;
int unamex ( Name)
struct xutsname *Name;
```

## Description

The **uname** subroutine stores information identifying the current system in the structure pointed to by the *Name* parameter.

The **uname** subroutine uses the **utsname** structure, which is defined in the **sys/utsname.h** file, and contains the following members:

```
char  sysname[SYS_NMLN];
char  nodename[SYS_NMLN];
char  release[SYS_NMLN];
char  version[SYS_NMLN];
char  machine[SYS_NMLN];
```

The **uname** subroutine returns a null-terminated character string naming the current system in the *sysname* character array. The *nodename* array contains the name that the system is known by on a communications network. The *release* and *version* arrays further identify the system. The *machine* array identifies the system unit hardware being used.

The **unamex** subroutine uses the **xutsname** structure, which is defined in the **sys/utsname.h** file, and contains the following members:

```
unsigned long  nid;
long  reserved[3];
```

The *xutsname.nid* field is the binary form of the *utsname.machine* field. For local area networks in which a binary node name is appropriate, the *xutsname.nid* field contains such a name.

Release and version variable numbers returned by the **uname** and **unamex** subroutines may change when new BOS software levels are installed. This change affects applications using these values to access licensed programs. Machine variable changes are due to hardware fixes or upgrades.

Contact the appropriate support organization if your application is affected.

## Parameters

*Name*     A pointer to the **utsname** or **xutsname** structure.

## Return Values

Upon successful completion, the **uname** or **unamex** subroutine returns a nonnegative value. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **uname** and **unamex** subroutines is unsuccessful if the following is true:

**EFAULT**     The *Name* parameter points outside of the process address space.

## Related Information

The **uname** command.

---

## ungetc or ungetwc Subroutine

### Purpose

Pushes a character back into the input stream.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int ungetc ( Character, Stream)
```

```
int Character;
```

```
FILE *Stream;
```

```
wint_t ungetwc (Character, Stream)
```

```
wint_t Character;
```

```
FILE *Stream;
```

### Description

The **ungetc** and **ungetwc** subroutines insert the character specified by the *Character* parameter (converted to an unsigned character in the case of the **ungetc** subroutine) into the buffer associated with the input stream specified by the *Stream* parameter. This causes the next call to the **getc** or **getwc** subroutine to return the *Character* value. A successful intervening call (with the stream specified by the *Stream* parameter) to a file-positioning subroutine (**fseek**, **fsetpos**, or **rewind**) discards any inserted characters for the stream. The **ungetc** and **ungetwc** subroutines return the *Character* value, and leaves the file (in its externally stored form) specified by the *Stream* parameter unchanged.

You can always push one character back onto a stream, provided that something has been read from the stream or the **setbuf** subroutine has been called. If the **ungetc** or **ungetwc** subroutine is called too many times on the same stream without an intervening read or file-positioning operation, the operation may not be successful. The **fseek** subroutine erases all memory of inserted characters.

The **ungetc** and **ungetwc** subroutines return a value of **EOF** or **WEOF** if a character cannot be inserted.

A successful call to the **ungetc** or **ungetwc** subroutine clears the end-of-file indicator for the stream specified by the *Stream* parameter. The value of the file-position indicator after all inserted characters are read or discarded is the same as before the characters were inserted. The value of the file-position indicator is decreased after each successful call to the **ungetc** or **ungetwc** subroutine. If its value was 0 before the call, its value is indeterminate after the call.

### Parameters

*Character*            Specifies a character.

*Stream*                Specifies the input stream.

### Return Values

The **ungetc** and **ungetwc** subroutines return the inserted character if successful; otherwise, **EOF** or **WEOF** is returned, respectively.

## Related Information

Other wide character I/O subroutines: **fgetwc** subroutine, **fgetws** subroutine, **fputwc** subroutine, **fputws** subroutine, **getwc** subroutine, **getwchar** subroutine, **getws** subroutine, **putwc** subroutine, **putwchar** subroutine, **putws** subroutine.

Related standard I/O subroutines: **fdopen** subroutine, **fgets** subroutine, **fopen** subroutine, **fprintf** subroutine, **fputc** subroutine, **fputs** subroutine, **fread** subroutine, **freopen** subroutine, **fwrite** subroutine, **gets** subroutine, **printf** subroutine, **putc** subroutine, **putchar** subroutine, **puts** subroutine, **putw** subroutine, **sprintf** subroutine.

Subroutines, Example Programs, and Libraries, in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Multibyte Code and Wide Character Code Conversion Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## unlink Subroutine

### Purpose

Removes a directory entry.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int unlink ( Path )
const char *Path;
```

### Description

The **unlink** subroutine removes the directory entry specified by the *Path* parameter and decreases the link count of the file referenced by the link. If Network File System (NFS) is installed on your system, this path can cross into another node.

**Attention:** Removing a link to a directory requires root user authority. Unlinking of directories is strongly discouraged since erroneous directory structures can result. The **rmdir** subroutine should be used to remove empty directories.

When all links to a file are removed and no process has the file open, all resources associated with the file are reclaimed, and the file is no longer accessible. If one or more processes have the file open when the last link is removed, the directory entry disappears. However, the removal of the file contents is postponed until all references to the file are closed.

If the parent directory of *Path* has the **sticky** attribute (described in the **mode.h** file), the calling process must have root user authority or an effective user ID equal to the owner ID of *Path* or the owner ID of the parent directory of *Path*.

The `st_ctime` and `st_mtime` fields of the parent directory are marked for update if the **unlink** subroutine is successful. In addition, if the file's link count is not 0, the `st_ctime` field of the file will be marked for update.

Applications should use the **rmdir** subroutine to remove a directory. If the *Path* parameter names a symbolic link, the link itself is removed.

## Parameters

*Path* Specifies the directory entry to be removed.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned, the **errno** global variable is set to indicate the error, and the specified file is not changed.

## Error Codes

The **unlink** subroutine fails and the named file is not unlinked if one of the following is true:

**ENOENT** The named file does not exist.

**EACCES** Write permission is denied on the directory containing the link to be removed.

**EBUSY** The entry to be unlinked is the mount point for a mounted filesystem, or the file named by *Path* is a named STREAM.

**EPERM** The file specified by the *Path* parameter is a directory, and the calling process does not have root user authority.

**EPERM** is also returned if the file named by the *Path* parameter is a directory in a JFS2 file system. Note that JFS allows you to unlink a directory.

**EROFS** The entry to be unlinked is part of a read-only file system.

The **unlink** subroutine can be unsuccessful for other reasons. For a list of additional errors, see Appendix A, "Base Operating System Error Codes for Service That Require Path-Name Resolution"

If NFS is installed on the system, the **unlink** subroutine can also fail if the following is true:

**ETIMEDOUT** The connection timed out.

## Related Information

The **close** subroutine, **link** subroutine, **open** subroutine, **remove** ("remove Subroutine" on page 40) subroutine, **rmdir** ("rmdir Subroutine" on page 48) subroutine.

The **rm** command.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## unload Subroutine

### Purpose

Unloads a module.

### Library

Standard C Library (**libc.a**)

## Syntax

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

```
int unload( FunctionPointer)  
int (*FunctionPointer)( );
```

## Description

The **unload** subroutine unloads the specified module and its dependents. The value returned by the **load** subroutine is passed to the **unload** subroutine as *FunctionPointer*. The **unload** subroutine calls termination routines (fini routines) for the specified module and any of its dependents that are not being used by any other module.

The **unload** subroutine frees the storage used by the specified module only if the module is no longer in use. A module is in use as long as any other module that is in use imports symbols from it.

When a module is unloaded, any deferred resolution symbols that were bound to the module remain bound. These bindings create references to the module that cannot be undone, even with the **unload** subroutine.

(This paragraph only applies to AIX 4.3.1 and previous releases.) When a process is executing under **ptrace** control, portions of the process's address space are recopied after the **unload** processing completes. For a 32-bit process, the main program text (loaded in segment 1) and shared library modules (loaded in segment 13) are recopied. Any breakpoints or other modifications to these segments must be reinserted after the **unload** call. For a 64-bit process, shared library modules are recopied after an **unload** call. The debugger will be notified by setting the **W\_SLWTED** flag in the status returned by **wait**, so that it can reinsert breakpoints.

(This paragraph only applies to AIX 4.3.2 and later releases.) When a process executing under **ptrace** control calls **unload**, the debugger is notified by setting the **W\_SLWTED** flag in the status returned by **wait**. If a module loaded in the shared library is no longer in use by the process, the module is deleted from the process's copy of the shared library segment by freeing the pages containing the module.

## Parameters

*FunctionPointer*                      Specifies the name of the function returned by the **load** subroutine.

## Return Values

Upon successful completion, the **unload** subroutine returns a value of 0, even if the module couldn't be unloaded because it is still in use.

## Error Codes

If the **unload** subroutine fails, a value of -1 is returned, the program is not unloaded, and **errno** is set to indicate the error. **errno** may be set to one of the following:

**EINVAL**     The *FunctionPointer* parameter does not correspond to a program loaded by the **load** subroutine.

## Related Information

The **load** subroutine, **loadbind** subroutine, **loadquery** subroutine, **dlclose** subroutine.

The **ld** command.

## unlockpt Subroutine

### Purpose

Unlocks a pseudo-terminal device.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int unlockpt ( FileDescriptor )  
int FileDescriptor;
```

### Description

The **unlockpt** subroutine unlocks the slave pseudo-terminal device associated with the master pseudo-terminal device defined by the *FileDescriptor* parameter. This subroutine has no effect if the environment variable XPG\_SUS\_ENV is not set equal to the string "ON", or if the BSD PTY driver is used.

### Parameters

*FileDescriptor*                      Specifies the file descriptor of the master pseudo-terminal device.

### Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

### Related Information

The **grantpt** subroutine.

The Input and Output Handling Programmer's Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## usrinfo Subroutine

### Purpose

Gets and sets user information about the owner of the current process.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <uinfo.h>
```

```
int usrinfo ( Command, Buffer, Count)
int Command;
char *Buffer;
int Count;
```

## Description

The **usrinfo** subroutine gets and sets information about the owner of the current process. The information is a sequence of null-terminated *name=value* strings. The last string in the sequence is terminated by two successive null characters. A child process inherits the user information of the parent process.

## Parameters

*Command* Specifies one of the following constants:

### GETUINFO

Copies user information, up to the number of bytes specified by the *Count* parameter, into the buffer pointed to by the *Buffer* parameter.

### SETUINFO

Sets the user information for the process to the number of bytes specified by the *Count* parameter in the buffer pointed to by the *Buffer* parameter. The calling process must have root user authority to set the user information.

The minimum user information consists of four strings typically set by the **login** program:

**NAME**=*UserName*

**LOGIN**=*LoginName*

**LOGNAME**=*LoginName*

**TTY**=*TTYName*

If the process has no terminal, the *TTYName* parameter should be null.

*Buffer* Specifies a pointer to a user buffer. This buffer is usually **UINFOSIZ** bytes long.

*Count* Specifies the number of bytes of user information copied from or to the user buffer.

## Return Values

If successful, the **usrinfo** subroutine returns a non-negative integer giving the number of bytes transferred. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **usrinfo** subroutine fails if one of the following is true:

- EPERM** The *Command* parameter is set to **SETUINFO**, and the calling process does not have root user authority.
- EINVAL** The *Command* parameter is not set to **SETUINFO** or **GETUINFO**.
- EINVAL** The *Command* parameter is set to **SETUINFO**, and the *Count* parameter is larger than **UINFOSIZ**.
- EFAULT** The *Buffer* parameter points outside of the address space of the process.

## Related Information

The **getuinfo** subroutine, **setpenv** (“setpenv Subroutine” on page 142) subroutine.

The **login** command.

List of Security and Auditing Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

## utimes or utime Subroutine

### Purpose

Sets file-access and modification times.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int utimes ( Path, Times)
```

```
char *Path;
```

```
struct timeval Times[2];
```

```
#include <utime.h>
```

```
int utime ( Path, Times)
```

```
const char *Path;
```

```
const struct utimbuf *Times;
```

### Description

The **utimes** subroutine sets the access and modification times of the file pointed to by the *Path* parameter to the value of the *Times* parameter. This subroutine allows time specifications accurate to the second.

The **utime** subroutine also sets file access and modification times. Each time is contained in a single integer and is accurate only to the nearest second. If successful, the **utime** subroutine marks the time of the last file-status change (**st\_ctime**) to be updated.

Microsecond time stamps are not implemented, even though the **utimes** subroutine provides a way to specify them.

### Parameters

*Path*      Points to the file.

*Times* Specifies the date and time of last access and of last modification. For the **utimes** subroutine, this is an array of **timeval** structures, as defined in the **sys/time.h** file. The first array element represents the date and time of last access, and the second element represents the date and time of last modification. The times in the **timeval** structure are measured in seconds and microseconds since 00:00:00 Greenwich Mean Time (GMT), 1 January 1970, rounded to the nearest second.

For the **utime** subroutine, this parameter is a pointer to a **utimbuf** structure, as defined in the **utime.h** file. The first structure member represents the date and time of last access, and the second member represents the date and time of last modification. The times in the **utimbuf** structure are measured in seconds since 00:00:00 Greenwich Mean Time (GMT), 1 January 1970.

If the *Times* parameter has a null value, the access and modification times of the file are set to the current time. If the file is remote, the current time at the remote node, rather than the local node, is used. To use the call this way, the effective user ID of the process must be the same as the owner of the file or must have root authority, or the process must have write permission to the file.

If the *Times* parameter does not have a null value, the access and modification times are set to the values contained in the designated structure, regardless of whether those times are the same as the current time. Only the owner of the file or a user with root authority can use the call this way.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned, the **errno** global variable is set to indicate the error, and the file times are not changed.

## Error Codes

The **utimes** or **utime** subroutine fails if one of the following is true:

- EPERM** The *Times* parameter is not null and the calling process neither owns the file nor has root user authority.
- EACCESS** The *Times* parameter is null, effective user ID is neither the owner of the file nor has root authority, or write access is denied.
- EROFS** The file system that contains the file is mounted read-only.

The **utimes** or **utime** subroutine can be unsuccessful for other reasons. For a list of additional errors, see "Base Operating System Error Codes For Services That Require Path-Name Resolution."

The **utimes** or **utime** subroutine can be unsuccessful for other reasons. For a list of additional errors, see Appendix A, "Base Operating System Error Codes For Services That Require Path-Name Resolution."

## Related Information

The **stat** ("statx, stat, lstat, fstatx, fstat, fullstat, fullstat, stat64, lstat64, or fstat64 Subroutine" on page 277) subroutine.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## varargs Macros

### Purpose

Handles a variable-length parameter list.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <stdarg.h>
```

```
type va_arg ( Argp, Type)  
va_list Argp;
```

```
void va_start (Argp, ParmN)  
va_list Argp;
```

```
void va_end (Argp)  
va_list Argp;
```

OR

```
#include <varargs.h>
```

```
va_alist Argp;  
va_dcl
```

```
void va_start (Argp)  
va_list Argp;
```

```
type va_arg (Argp, Type)  
va_list Argp;
```

```
void va_end (Argp)  
va_list Argp;
```

## Description

The **varargs** set of macros allows you to write portable subroutines that accept a variable number of parameters. Subroutines that have variable-length parameter lists (such as the **printf** subroutine), but that do not use the **varargs** macros, are inherently nonportable because different systems use different parameter-passing conventions.

**Note:** Do not include both **<stdarg.h>** and **<varargs.h>**. Use of **<varargs.h>** is not recommended. It is supplied for backwards compatibility.

### For <stdarg.h>

**va\_start**            Initializes the *Argp* parameter to point to the beginning of the list. The *ParmN* parameter identifies the rightmost parameter in the function definition. For compatibility with previous programs, it defaults to the address of the first parameter on the parameter list. Acceptable parameters include: integer, double, and pointer. The **va\_start** macro is started before any access to the unnamed arguments.

### For <varargs.h>

**va\_alist**            A variable used as the parameter list in the function header.  
**va\_argp**             A variable that the **varargs** macros use to keep track of the current location in the parameter list. Do not modify this variable.  
**va\_dcl**              Declaration for **va\_alist**. No semicolon should follow **va\_dcl**.  
**va\_start**             Initializes the *Argp* parameter to point to the beginning of the list.

### For <stdarg.h> and <varargs.h>

**va\_list**             Defines the type of the variable used to traverse the list.

**va\_arg** Returns the next parameter in the list pointed to by the *Argp* parameter.  
**va\_end** Cleans up at the end.

Your subroutine can traverse, or scan, the parameter list more than once. Start each traversal with a call to the **va\_start** macro and end it with the **va\_end** macro.

**Note:** The calling routine is responsible for specifying the number of parameters because it is not always possible to determine this from the stack frame. For example, **execl** is passed a null pointer to signal the end of the list. The **printf** subroutine determines the number of parameters from its *Format* parameter.

## Parameters

*Argp* Specifies a variable that the **varargs** macros use to keep track of the current location in the parameter list. Do not modify this variable.  
*Type* Specifies the type to which the expected argument will be converted when passed as an argument. In C, arguments that are char or short should be accessed as int; unsigned char or short arguments are converted to unsigned int, and float arguments are converted to double. Different types can be mixed, but it is up to the routine to know what type of argument is expected, because it cannot be determined at runtime.  
*ParmN* Specifies a parameter that is the identifier of the rightmost parameter in the function definition.

## Examples

The following **execl** system call implementations are examples of the **varargs** macros usage.

1. The following example includes **<stdarg.h>**:

```
#include <stdarg.h>
#define MAXargs 31
int execl (const char *path, ...)
{
    va_list Argp;
    char *array [MAXargs];
    int argno=0;
    va_start (Argp, path);
    while ((array[argno++] = va_arg(Argp, char*)) != (char*)0)
        ;
    va_end(Argp);
    return(execv(path, array));
}
main()
{
    execl("/usr/bin/echo", "ArgV[0]", "This", "Is", "A", "Test", "\0");
    /* ArguementV[0] will be discarded by the execv in main(): */
    /* by convention ArgV[0] should be a copy of path parameter */
}
```

2. The following example includes **<varargs.h>**:

```
#include <varargs.h>
#define MAXargS 100
/*
** execl is called by
** execl(file, arg1, arg2, . . . , (char *) 0);
*/
execl(va_alist)
    va_dcl
{
    va_list ap;
    char *file;
    char *args[MAXargS];
    int argno = 0;
```

```

    va_start(ap);
    file = va_arg(ap, char *);
    while ((args[argno++] = va_arg(ap, char *)) != (char *) 0)
        ; /* Empty loop body */
    va_end(ap);
    return (execv(file, args));
}

```

## Related Information

The **exec** subroutines.

The **printf** subroutine.

List of String Manipulation Services in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## vfscanf, vscanf, or vsscanf Subroutine

### Purpose

Formats input of an argument list.

### Syntax

```

#include <stdarg.h>
#include <stdio.h>

```

```

int vfscanf (stream, format, arg)
File *restrict stream
const char format;
va_list arg;

```

```

int vscanf (format, arg)
const char format;
va_list arg;

```

```

int vsscanf (format, arg)
const char format;
va_list arg;

```

### Description

The **vscanf**, **vfscanf**, and **vsscanf** subroutines are equivalent to the **scanf**, **fscanf**, and **sscanf** subroutines, 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 file. These subroutines do not invoke the **va\_end** macro. As these functions invoke the **va\_arg** macro, the value of *ap* after the return is unspecified.

### Parameters

*stream*  
*format*  
*arg*

### Return Values

Upon successful completion, these functions shall return the number of successfully matched and assigned input items; this number can be zero in the event of an early matching failure. If the input ends

before the first matching failure or conversion, EOF shall be returned. If a read error occurs, the error indicator for the stream is set, EOF shall be returned, and *errno* shall be set to indicate the error.

## Related Information

The “scanf, fscanf, sscanf, or wscanf Subroutine” on page 109.

---

## vwscanf, vswscanf, or vwscanf Subroutine

### Purpose

Wide-character formatted input of the argument list.

### Syntax

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

```
int vwscanf (stream, format, arg)
FILE *restrict stream;
const wchar_t format;
va_list arg;
```

```
int vswscanf (ws, format, arg)
const wchar_t *restrict ws;
const wchar_t format;
va_list arg;
```

```
int vwscanf (format, arg)
const wchar_t format;
va_list arg;
```

### Description

The **vwscanf**, **vswscanf**, and **vwscanf** subroutines are equivalent to the **fscanf**, **swscanf**, and **wscanf** subroutines, 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 file. These subroutines do not invoke the **va\_end** macro. As these subroutines invoke the **va\_arg** macro, the value of *ap* after the return is unspecified.

### Return Values

Upon successful completion, the **vwscanf**, **vswscanf**, and **vwscanf** subroutines return the number of successfully matched and assigned input items. This number can be zero 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 the **errno** global variable is set to indicate the error.

### Related Information

fscanf, wscanf, swscanf Subroutines in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

---

## vwprintf, vwprintf Subroutine

### Purpose

Wide-character formatted output of a stdarg argument list.

## Library

Standard library (**libc.a**)

## Syntax

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

int vwprintf ((const wchar_t * format, va_list arg) );
int vfwprintf(FILE * stream, const wchar_t * format, va_list arg);
int vswprintf (wchar_t * s, size_t n, const wchar_t * format, va_list arg);
```

## Description

The **vwprintf**, **vfwprintf** and **vsprintf** 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**.

These functions do not invoke the **va\_end** macro. However, as these functions do invoke the **va\_arg** macro, the value of **ap** after the return is indeterminate.

## Return Values

Refer to **fwprintf**.

## Error Codes

Refer to **fwprintf**.

## Related Information

The **fwprintf** subroutine.

---

## vmgetinfo Subroutine

### Purpose

Retrieves Virtual Memory Manager (VMM) information.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int vmgetinfo(void *out, int command, int arg)
```

### Description

The **vmgetinfo** subroutine returns the current value of certain VMM parameters.

### Parameters

*arg* Additional parameter which depends on the command parameter.

*command* Specifies which information should be returned. The *command* parameter has the following valid value:

**VMINFO**  
The content of the **vminfo** structure (described in **sys/vminfo.h**) is returned. The *out* parameter should point to a **vminfo** structure and *arg* should be the size of this structure. The smaller of the *arg* or *sizeof* (struct **vminfo**) parameters will be copied.

**VM\_PAGE\_INFO**  
The size, in bytes, of the page backing the address specified in the *addr* field of the **vm\_page\_info** structure (described in **sys/vminfo.h**) is returned. The *out* parameter should point to a **vm\_page\_info** structure with the *addr* field set to the desired address of which to query the page size. The *arg* parameter should be the size of the **vm\_page\_info** structure.

**IPC\_LIMITS**  
The content of the **ipc\_limits** struct (described in the **sys/vminfo.h** file) is returned. The *out* parameter should point to an **ipc\_limits** structure and *arg* should be the size of this structure. The smaller of the *arg* or *sizeof* (struct **ipc\_limits**) parameters will be copied. The **ipc\_limits** struct contains the inter-process communication (IPC) limits for the system.

*out* Specifies the address where VMM information should be returned.

## Return Values

If the **vmgetinfo** subroutine is successful, a value of 0 is returned. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **vmgetinfo** does not succeed if the following are true:

<b>EFAULT</b>	The copy operation to the buffer was not successful.
<b>ENOSYS</b>	The command parameter is not valid (or not yet implemented).
<b>EINVAL</b>	When VM_PAGE_INFO is the command, the <i>addr</i> field of the <b>vm_page_info</b> structure is an invalid address.

---

## vmount or mount Subroutine

### Purpose

Makes a file system available for use.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int vmount ( VMount, Size)
struct vmount *VMount;
int Size;
```

```
int mount
( Device, Path, Flags)
```

```
char *Device;  
char *Path;  
int  Flags;
```

## Description

The **vmount** subroutine mounts a file system, thereby making the file available for use. The **vmount** subroutine effectively creates what is known as a *virtual file system*. After a file system is mounted, references to the path name that is to be mounted over refer to the root directory on the mounted file system.

A directory can only be mounted over a directory, and a file can only be mounted over a file. (The file or directory may be a symbolic link.)

Therefore, the **vmount** subroutine can provide the following types of mounts:

- A local file over a local or remote file
- A local directory over a local or remote directory
- A remote file over a local or remote file
- A remote directory over a local or remote directory.

A mount to a directory or a file can be issued if the calling process has root user authority or is in the system group and has write access to the mount point.

To mount a block device, remote file, or remote directory, the calling process must also have root user authority.

The **mount** subroutine only allows mounts of a block device over a local directory with the default file system type. The **mount** subroutine searches the **/etc/filesystems** file to find a corresponding stanza for the desired file system.

**Note:** The **mount** subroutine interface is provided only for compatibility with previous releases of the operating system. The use of the **mount** subroutine is strongly discouraged by normal application programs.

If the directory you are trying to mount over has the sticky bit set to on, you must either own that directory or be the root user for the mount to succeed. This restriction applies only to directory-over-directory mounts.

## Parameters

### *Device*

A path name identifying the block device (also called a special file) that contains the physical file system.

### *Path*

A path name identifying the directory on which the file system is to be mounted.

### *Flags*

Values that define characteristics of the object to be mounted. Currently these values are defined in the **/usr/include/sys/vmount.h** file:

#### **MNT\_READONLY**

Indicates that the object to be mounted is read-only and that write access is not allowed. If this value is not specified, writing is permitted according to individual file accessibility.

#### **MNT\_NOSUID**

Indicates that **setuid** and **setgid** programs referenced through the mount should not be executable. If this value is not specified, **setuid** and **setgid** programs referenced through the mount may be executable.

## **MNT\_NODEV**

Indicates that opens of device special files referenced through the mount should not succeed. If this value is not specified, opens of device special files referenced through the mount may succeed.

## *VMount*

A pointer to a variable-length **vmount** structure. This structure is defined in the **sys/vmount.h** file.

The following fields of the *VMount* parameter must be initialized before the call to the **vmount** subroutine:

### **vmt\_revision**

The revision code in effect when the program that created this virtual file system was compiled. This is the value **VMT\_REVISION**.

### **vmt\_length**

The total length of the structure with all its data. This must be a multiple of the word size (4 bytes) and correspond with the *Size* parameter.

### **vmt\_flags**

Contains the general mount characteristics. The following value may be specified:

#### **MNT\_READONLY**

A read-only virtual file system is to be created.

### **vmt\_gfstype**

The type of the generic file system underlying the **VMT\_OBJECT**. Values for this field are defined in the **sys/vmount.h** file and include:

#### **MNT\_JFS**

Indicates the native file system.

#### **MNT\_NFS**

Indicates a Network File System client.

#### **MNT\_CDROM**

Indicates a CD-ROM file system.

### **vmt\_data**

An array of structures that describe variable length data associated with the **vmount** structure. The structure consists of the following fields:

#### **vmt\_off**

The offset of the data from the beginning of the **vmount** structure.

### **vmt\_size**

The size, in bytes, of the data.

The array consists of the following fields:

#### **vmt\_data[VMT\_OBJECT]**

Specifies the name of the device, directory, or file to be mounted.

#### **vmt\_data[VMT\_STUB]**

Specifies the name of the device, directory, or file to be mounted over.

#### **vmt\_data[VMT\_HOST]**

Specifies the short (binary) name of the host that owns the mounted object. This need not be specified if **VMT\_OBJECT** is local (that is, it has the same *vmt\_gfstype* as / (root), the root of all file systems).

#### **vmt\_data[VMT\_HOSTNAME]**

Specifies the long (character) name of the host that owns the mounted object. This need not be specified if **VMT\_OBJECT** is local.

**vmt\_data[VMT\_INFO]**

Specifies binary information to be passed to the generic file-system implementation that supports **VMT\_OBJECT**. The interpretation of this field is specific to the `gfs_type`.

**vmt\_data[VMT\_ARGS]**

Specifies a character string representation of **VMT\_INFO**.

On return from the **vmount** subroutine, the following additional fields of the *VMount* parameter are initialized:

**vmt\_fsid**

Specifies the two-word file system identifier; the interpretation of this identifier depends on the `gfs_type`.

**vmt\_vfsnumber**

Specifies the unique identifier of the virtual file system. Virtual file systems do not survive the IPL; neither does this identifier.

**vmt\_time**

Specifies the time at which the virtual file system was created.

*Size* Specifies the size, in bytes, of the supplied data area.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned, and the **errno** global variable is set to indicate the error.

## Error Codes

The **mount** and **vmount** subroutines fail and the virtual file system is not created if any of the following is true:

<b>EACCES</b>	The calling process does not have write permission on the stub directory (the directory to be mounted over).
<b>EBUSY</b>	<b>VMT_OBJECT</b> specifies a device that is already mounted or an object that is open for writing, or the kernel's mount table is full.
<b>EFAULT</b>	The <i>VMount</i> parameter points to a location outside of the allocated address space of the process.
<b>EFBIG</b>	The size of the file system is too big.
<b>EFORMAT</b>	An internal inconsistency has been detected in the file system.
<b>EINVAL</b>	The contents of the <i>VMount</i> parameter are unintelligible (for example, the <code>vmt_gfstype</code> is unrecognizable, or the file system implementation does not understand the <b>VMT_INFO</b> provided).
<b>ENOSYS</b>	The file system type requested has not been configured.
<b>ENOTBLK</b>	The object to be mounted is not a file, directory, or device.
<b>ENOTDIR</b>	The types of <b>VMT_OBJECT</b> and <b>VMT_STUB</b> are incompatible.
<b>EPERM</b>	<b>VMT_OBJECT</b> specifies a block device, and the calling process does not have root user authority.
<b>EROFS</b>	An attempt has been made to mount a file system for read/write when the file system cannot support writing.

The **mount** and **vmount** subroutines can also fail if additional errors occur.

## Related Information

The **mntctl** subroutine, **umount** (“umount or uvmount Subroutine” on page 419) subroutine.

The **mount** command, **umount** command.

Files, Directories, and File Systems for Programmers in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## vsnprintf Subroutine

### Purpose

Print formatted output.

### Library

Standard library (**libc.a**)

### Syntax

```
#include <stdarg.h>
#include <stdio.h>
int vsnprintf(char * s, size_t n, const char * format, va_list ap)
```

### Description

Refer to **fprintf**.

---

## vwsprintf Subroutine

### Purpose

Writes formatted wide characters.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <wchar.h>
#include <stdarg.h>

int vwsprintf (wcs, Format, arg)
wchar_t * wcs;
const char * Format;
va_list arg;
```

### Description

The **vwsprintf** subroutine writes formatted wide characters. It is structured like the **vsprintf** subroutine with a few differences. One difference is that the *wcs* parameter specifies a wide character array into which the generated output is to be written, rather than a character array. The second difference is that the meaning of the **S** conversion specifier is always the same in the case where the **#** flag is specified. If copying takes place between objects that overlap, the behavior is undefined.

### Parameters

<i>wcs</i>	Specifies the array of wide characters where the output is to be written.
<i>Format</i>	Specifies a multibyte character sequence composed of zero or more directives (ordinary multibyte characters and conversion specifiers). The new formats added to handle the wide characters are:  <b>%C</b> Formats a single wide character.  <b>%S</b> Formats a wide character string.
<i>arg</i>	Specifies the parameters to be printed.

## Return Values

The **wvsprintf** subroutine returns the number of wide characters (not including the terminating wide character null) written into the wide character array and specified by the *wcs* parameter.

## Related Information

The **vsprintf** subroutine.

The **printf** command.

National Language Support Overview in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## wait, waitpid, wait3, or wait364 Subroutine

### Purpose

Waits for a child process to stop or terminate.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/wait.h>
pid_t wait ( StatusLocation)
int *StatusLocation;
pid_t wait ((void *) 0)
#include <sys/wait.h>
```

```
pid_t waitpid ( ProcessID,
StatusLocation, Options)
int *StatusLocation;
pid_t ProcessID;
int Options;
#include <sys/time.h>
#include <sys/resource.h>
#include <sys/wait.h>
```

```
pid_t wait3 (StatusLocation, Options, ResourceUsage)
int *StatusLocation;
int Options;
struct rusage *ResourceUsage;
```

```
pid_t wait364 (StatusLocation, Options, ResourceUsage)
int *StatusLocation;
int Options;
struct rusage64 *ResourceUsage;
```

### Description

The **wait** subroutine suspends the calling thread until the process receives a signal that is not blocked or ignored, or until any one of the calling process' child processes stops or terminates. The **wait** subroutine returns without waiting if the child process that has not been waited for has already stopped or terminated prior to the call.

**Note:** The effect of the **wait** subroutine can be modified by the setting of the **SIGCHLD** signal. See the **sigaction** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine for details.

The **waitpid** subroutine includes a *ProcessID* parameter that allows the calling thread to gather status from a specific set of child processes, according to the following rules:

- If the *ProcessID* value is equal to a value of -1, status is requested for any child process. In this respect, the **waitpid** subroutine is equivalent to the **wait** subroutine.
- A *ProcessID* value that is greater than 0 specifies the process ID of a single child process for which status is requested.
- If the *ProcessID* parameter is equal to 0, status is requested for any child process whose process group ID is equal to that of the calling thread’s process.
- If the *ProcessID* parameter is less than 0, status is requested for any child process whose process group ID is equal to the absolute value of the *ProcessID* parameter.

The **waitpid**, **wait3**, and **wait364** subroutine variants provide an *Options* parameter that can modify the behavior of the subroutine. Two values are defined, **WNOHANG** and **WUNTRACED**, which can be combined by specifying their bitwise-inclusive OR. The **WNOHANG** option prevents the calling thread from being suspended even if there are child processes to wait for. In this case, a value of 0 is returned indicating there are no child processes that have stopped or terminated. If the **WUNTRACED** option is set, the call should also return information when children of the current process are stopped because they received a **SIGTTIN**, **SIGTTOU**, **SIGSSTP**, or **SIGTSTOP** signal.

The **wait364** subroutine can be called to make 64-bit *rusage* counters explicitly available in a 32-bit environment.

In AIX 5.1 and later, 64-bit quantities are also available to 64-bit applications through the **wait3()** interface in the *ru\_utime* and *ru\_stime* fields of **struct rusage**.

When a 32-bit process is being debugged with **ptrace**, the status location is set to **W\_SLWTED** if the process calls **load**, **unload**, or **loadbind**. When a 64-bit process is being debugged with **ptrace**, the status location is set to **W\_SLWTED** if the process calls **load** or **unload**.

If multiprocessing debugging mode is enabled, the status location is set to **W\_SEWTED** if a process is stopped during an **exec** subroutine and to **W\_SFWTED** if the process is stopped during a **fork** subroutine.

If more than one thread is suspended awaiting termination of the same child process, exactly one thread returns the process status at the time of the child process termination.

If the **WCONTINUED** option is set, the call should return information when the children of the current process have been continued from a job control stop but whose status has not yet been reported.

## Parameters

<i>StatusLocation</i>	Points to integer variable that contains (or will contain) the child process termination status, as defined in the <b>sys/wait.h</b> file.
<i>ProcessID</i>	Specifies the child process.
<i>Options</i>	Modifies behavior of subroutine.
<i>ResourceUsage</i>	Specifies the location of a structure to be filled in with resource utilization information for terminated children.

## Macros

The value pointed to by *StatusLocation* when **wait**, **waitpid**, or **wait3** subroutines are returned, can be used as the *ReturnedValue* parameter for the following macros defined in the **sys/wait.h** file to get more information about the process and its child process.

**WIFCONTINUED**(*ReturnedValue*)  
**pid\_t** *ReturnedValue*;

Returns a nonzero value if status returned for a child process that has continued from a job control stop.

**WIFSTOPPED**(*ReturnedValue*)  
**int** *ReturnedValue*;

Returns a nonzero value if status returned for a stopped child.

**int**  
**WSTOPSIG**(*ReturnedValue*)  
**int** *ReturnedValue*;

Returns the number of the signal that caused the child to stop.

**WIFEXITED**(*ReturnedValue*)  
**int** *ReturnedValue*;

Returns a nonzero value if status returned for normal termination.

**int**  
**WEXITSTATUS**(*ReturnedValue*)  
**int** *ReturnedValue*;

Returns the low-order 8 bits of the child exit status.

**WIFSIGNALED**(*ReturnedValue*)  
**int** *ReturnedValue*;

Returns a nonzero value if status returned for abnormal termination.

**int**  
**WTERMSIG**(*ReturnedValue*)  
**int** *ReturnedValue*;

Returns the number of the signal that caused the child to terminate.

## Return Values

If the **wait** subroutine is unsuccessful, a value of -1 is returned and the **errno** global variable is set to indicate the error. In addition, the **waitpid**, **wait3**, and **wait364** subroutines return a value of 0 if there are no stopped or exited child processes, and the **WNOHANG** option was specified. The **wait** subroutine returns a 0 if there are no stopped or exited child processes, also.

## Error Codes

The **wait**, **waitpid**, **wait3**, and **wait364** subroutines are unsuccessful if one of the following is true:

<b>ECHILD</b>	The calling thread's process has no existing unwaited-for child processes.
<b>EINTR</b>	This subroutine was terminated by receipt of a signal.
<b>EFAULT</b>	The <i>StatusLocation</i> or <i>ResourceUsage</i> parameter points to a location outside of the address space of the process.

The **waitpid** subroutine is unsuccessful if the following is true:

<b>ECHILD</b>	The process or process group ID specified by the <i>ProcessID</i> parameter does not exist or is not a child process of the calling process.
---------------	--

The **waitpid** and **wait3** subroutines are unsuccessful if the following is true:

<b>EINVAL</b>	The value of the <i>Options</i> parameter is not valid.
---------------	---

## Related Information

The **exec** subroutine, **\_exit**, **exit**, or **atexit** subroutine, **fork** subroutine, **getrusage** subroutine, **pause** subroutine, **ptrace** subroutine, **sigaction** (“sigaction, sigvec, or signal Subroutine” on page 164) subroutine.

---

## waitid Subroutine

### Purpose

Waits for a child process to change state.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/wait.h>;

int waitid (idtype, id, infop, options)
idtype_t idtype;
id_t id;
siginfo_t *infop;
int options;
```

### Description

The **waitid** subroutine suspends the calling thread until one child of the process containing the calling thread changes state. It records the current state of a child in the structure pointed to by the *infop* parameter. If a child process changed state prior to the call to the **waitid** subroutine, the **waitid** subroutine returns immediately. If more than one thread is suspended in the **wait** or **waitpid** subroutines waiting for termination of the same process, exactly one thread will return the process status at the time of the target process termination.

### Parameters

<i>idtype</i>	Specifies the child process.
<i>id</i>	Specifies the child process.
<i>infop</i>	Specifies the location of a <b>siginfo_t</b> structure to be filled in with resource utilization information.
<i>options</i>	Specifies which state changes the <b>waitid</b> subroutine will wait for. It is formed by OR'ing together one or more of the following flags:
	<b>WEXITED</b> Wait for processes that have exited.
	<b>WSTOPPED</b> Status will be returned for any child that has stopped upon receipt of a signal.
	<b>WCONTINUED</b> Status will be returned for any child that was stopped and has been continued.
	<b>WNOHANG</b> Return immediately if there are no children to wait for.
	<b>WNOWAIT</b> Keep the process whose status is returned in the <i>infop</i> parameter in a waitable state. This will not affect the state of the process. The process can be waited for again after this call completes.

## Return Values

If **WNOHANG** was specified and there are no children to wait for, 0 is returned. If the **waitid** subroutine returns due to the change of state of one of its children, 0 is returned. Otherwise, -1 is returned and **errno** is set to indicate the error.

## Error Codes

The **waitid** subroutine will fail if:

**ECHILD**

The calling process has no existing unwaited-for child processes.

**EINTR**

The **waitid** subroutine was interrupted by a signal.

**EINVAL**

An invalid value was specified for the *options*, or *idtype* parameters and the *id* parameter specifies an invalid set of processes.

## Related Information

The “wait, waitpid, wait3, or wait364 Subroutine” on page 440.

The `exec` subroutine and `_exit`, `exit`, or `atexit` subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

---

## wcscat, wcschr, wcscmp, wcsncpy, or wcsncpy Subroutine

### Purpose

Performs operations on wide-character strings.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
wchar_t * wcscat(WcString1, WcString2)
wchar_t * WcString1;
const wchar_t * WcString2;
```

```
wchar_t * wcschr(WcString, WideCharacter)
const wchar_t *WcString;
wchar_t WideCharacter;
```

```
int * wcscmp (WcString1, WcString2)
const wchar_t *WcString1, *WcString2;
```

```
wchar_t * wcsncpy(WcString1, WcString2)
wchar_t *WcString1;
const wchar_t
*
WcString2;
```

```
size_t wcsncpy(WcString1, WcString2)
const wchar_t *WcString1, *WcString2;
```

## Description

The **wcscat**, **wcschr**, **wscmp**, **wscpy**, or **wscspn** subroutine operates on null-terminated **wchar\_t** strings. These subroutines expect the string arguments to contain a **wchar\_t** null character marking the end of the string. A copy or concatenation operation does not perform boundary checking.

The **wcscat** subroutine appends a copy of the wide-character string pointed to by the *WcString2* parameter (including the terminating null wide-character code) to the end of the wide-character string pointed to by the *WcString1* parameter. The initial wide-character code of the *WcString2* parameter overwrites the null wide-character code at the end of the *WcString1* parameter. If successful, the **wcscat** subroutine returns the *WcString1* parameter.

The **wcschr** subroutine returns a pointer to the first occurrence of the *WideCharacter* parameter in the *WcString* parameter. The character value may be a **wchar\_t** null character. The **wchar\_t** null character at the end of the string is included in the search. The **wcschr** subroutine returns a pointer to the wide character code, if found, or returns a null pointer if the wide character is not found.

The **wscmp** subroutine compares two **wchar\_t** strings. It returns an integer greater than 0 if the *WcString1* parameter is greater than the *WcString2* parameter. It returns 0 if the two strings are equivalent. It returns a number less than 0 if the *WcString1* parameter is less than the *WcString2* parameter. The sign of the difference in value between the first pair of wide-character codes that differ in the objects being compared determines the sign of a nonzero return value.

The **wscpy** subroutine copies the contents of the *WcString2* parameter (including the ending **wchar\_t** null character) into the *WcString1* parameter. If successful, the **wscpy** subroutine returns the *WcString1* parameter. If the **wscpy** subroutine copies between overlapping objects, the result is undefined.

The **wscspn** subroutine computes the number of **wchar\_t** characters in the initial segment of the string pointed to by the *WcString1* parameter that do not appear in the string pointed to by the *WcString2* parameter. If successful, the **wscspn** subroutine returns the number of **wchar\_t** characters in the segment.

## Parameters

<i>WcString1</i>	Points to a wide-character string.
<i>WcString2</i>	Points to a wide-character string.
<i>WideCharacter</i>	Specifies a wide character for location.

## Return Values

Upon successful completion, the **wcscat** and **wscpy** subroutines return a value of *ws1*. The **wcschr** subroutine returns a pointer to the wide character code. Otherwise, a null pointer is returned.

The **wscmp** subroutine returns an integer greater than, equal to, or less than 0, if the wide character string pointed to by the *WcString1* parameter is greater than, equal to, or less than the wide character string pointed to by the *WcString2* parameter.

The **wscspn** subroutine returns the length of the segment.

## Related Information

The **mbcat** subroutine, **mbchr** subroutine, **mbcmp** subroutine, **mbcpy** subroutine, **mbchr** subroutine, **wcsncat** (“*wcsncat*, *wcsncmp*, or *wcsncpy* Subroutine” on page 449) subroutine, **wcsncmp**

("wcsncat, wcsncmp, or wcsncpy Subroutine" on page 449) subroutine, **wcsncpy** ("wcsncat, wcsncmp, or wcsncpy Subroutine" on page 449) subroutine, **wcsrchr** ("wcsrchr Subroutine" on page 451) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview, Multibyte and Wide Character String Comparison Subroutines, Understanding Wide Character String Copy Subroutines, and Wide Character String Search Subroutine in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## wcscoll Subroutine

### Purpose

Compares wide character strings.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int wcscoll ( WcString1, WcString2)  
const wchar_t *WcString1, *WcString2;
```

### Description

The **wcscoll** subroutine compares the two wide-character strings pointed to by the *WcString1* and *WcString2* parameters based on the collation values specified by the **LC\_COLLATE** environment variable of the current locale.

**Note:** The **wcscoll** subroutine differs from the **wcsncmp** subroutine in that the **wcscoll** subroutine compares wide characters based on their collation values, while the **wcsncmp** subroutine compares wide characters based on their ordinal values. The **wcscoll** subroutine uses more time than the **wcsncmp** subroutine because it obtains the collation values from the current locale.

The **wcscoll** subroutine may be unsuccessful if the wide character strings specified by the *WcString1* or *WcString2* parameter contains characters outside the domain of the current collating sequence.

### Parameters

<i>WcString1</i>	Points to a wide-character string.
<i>WcString2</i>	Points to a wide-character string.

### Return Values

The **wcscoll** subroutine returns the following values:

< 0	The collation value of the <i>WcString1</i> parameter is less than that of the <i>WcString2</i> parameter.
= 0	The collation value of the <i>WcString1</i> parameter is equal to that of the <i>WcString2</i> parameter.
> 0	The collation value of the <i>WcString1</i> parameter is greater than that of the <i>WcString2</i> parameter.

The **wcscoll** subroutine indicates error conditions by setting the **errno** global variable. However, there is no return value to indicate an error. To check for errors, the **errno** global variable should be set to 0, then checked upon return from the **wcscoll** subroutine. If the **errno** global variable is nonzero, an error occurred.

## Error Codes

**EINVAL** The *WcString1* or *WcString2* arguments contain wide-character codes outside the domain of the collating sequence.

## Related Information

The **wscmp** (“*wscat, wcschr, wscmp, wcsncpy, or wcsncpy Subroutine*” on page 444) subroutine.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Understanding Wide Character String Collation Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## wcsftime Subroutine

### Purpose

Converts date and time into a wide character string.

### Library

Standard C Library (**libc. a**)

### Syntax

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

```
size_t wcsftime (WcString, Maxsize, Format, TimPtr)
wchar_t * WcString;
size_t Maxsize;
const wchar_t * Format;
const struct tm * TimPtr;
```

### Description

The **wcsftime** function is equivalent to the **strftime** function, except that:

- The argument *wcs* points to the initial element of an array of wide-characters into which the generated output is to be placed.
- The argument *maxsize* indicates the maximum number of wide-characters to be placed in the output array.
- The argument *format* 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.

If copying takes place between objects that overlap, the behavior is undefined.

### Parameters

*WcString* Contains the output of the **wcsftime** subroutine.

<i>Maxsize</i>	Specifies the maximum number of bytes (including the wide character null-terminating byte) that may be placed in the <i>WcString</i> parameter.
<i>Format</i>	Specifiers are the same as in <b>strftime</b> (“strftime Subroutine” on page 288) function.
<i>TimPtr</i>	Contains the data to be converted by the <b>wcsftime</b> subroutine.

## Return Values

If successful, and if the number of resulting wide characters (including the wide character null-terminating byte) is no more than the number of bytes specified by the *Maxsize* parameter, the **wcsftime** subroutine returns the number of wide characters (not including the wide character null-terminating byte) placed in the *WcString* parameter. Otherwise, 0 is returned and the contents of the *WcString* parameter are indeterminate.

## Related Information

The **mbstowcs** subroutine, **strfmon** (“strfmon Subroutine” on page 286) subroutine, **strftime** (“strftime Subroutine” on page 288) subroutine, **strptime** (“strptime Subroutine” on page 301) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and List of Time and Monetary Formatting Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

## wcsid Subroutine

### Purpose

Returns the charsetID of a wide character.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int wcsid ( WC)
const wchar_t WC;
```

### Description

The **wcsid** subroutine returns the charsetID of the **wchar\_t** character. No validation of the character is performed. The parameter must point to a value in the character range of the current code set defined in the current locale.

### Parameters

*WC* Specifies the character to be tested.

### Return Values

Successful completion returns an integer value representing the charsetID of the character. This integer can be a number from 0 through *n*, where *n* is the maximum character set defined in the CHARSETID field of the **charmap**. See “Understanding the Character Set Description (charmap) Source File” in *AIX 5L Version 5.2 National Language Support Guide and Reference* for more information.

## Related Information

The **csid** subroutine, **mbstowcs** subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

National Language Support Overview, Multibyte Code and Wide Character Code Conversion Subroutines, and Understanding the Character Set Description (charmap) Source File in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## wcslen Subroutine

### Purpose

Determines the number of characters in a wide-character string.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <wcstr.h>
```

```
size_t wcslen( WcString)  
const wchar_t *WcString;
```

### Description

The **wcslen** subroutine computes the number of **wchar\_t** characters in the string pointed to by the *WcString* parameter.

### Parameters

*WcString*        Specifies a wide-character string.

### Return Values

The **wcslen** subroutine returns the number of **wchar\_t** characters that precede the terminating **wchar\_t** null character.

## Related Information

The **mbslen** subroutine, **wctomb** (“wctomb Subroutine” on page 468) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Multibyte Code and Wide Character Code Conversion Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## wcsncat, wcsncmp, or wcsncpy Subroutine

### Purpose

Performs operations on a specified number of wide characters from one string to another.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <wchar.h>
```

```
wchar_t * wcsncat (WcString1, WcString2, Number)
wchar_t * WcString1;
const wchar_t * WcString2;
size_t Number;
```

```
wchar_t * wcsncmp (WcString1, WcString2, Number)
const wchar_t *WcString1, *WcString2;
size_t Number;
```

```
wchar_t * wcsncpy (WcString1, WcString2, Number)
wchar_t *WcString1;
const wchar_t *WcString2;
size_t Number;
```

## Description

The **wcsncat**, **wcsncmp** and **wcsncpy** subroutines operate on null-terminated wide character strings.

The **wcsncat** subroutine appends characters from the *WcString2* parameter, up to the value of the *Number* parameter, to the end of the *WcString1* parameter. It appends a **wchar\_t** null character to the result and returns the *WcString1* value.

The **wcsncmp** subroutine compares wide characters in the *WcString1* parameter, up to the value of the *Number* parameter, to the *WcString2* parameter. It returns an integer greater than 0 if the value of the *WcString1* parameter is greater than the value of the *WcString2* parameter. It returns a 0 if the strings are equivalent. It returns an integer less than 0 if the value of the *WcString1* parameter is less than the value of the *WcString2* parameter.

The **wcsncpy** subroutine copies wide characters from the *WcString2* parameter, up to the value of the *Number* parameter, to the *WcString1* parameter. It returns the value of the *WcString1* parameter. If the number of characters in the *WcString2* parameter is less than the *Number* parameter, the *WcString1* parameter is padded out with **wchar\_t null** characters to a number equal to the value of the *Number* parameter.

## Parameters

<i>WcString1</i>	Specifies a wide-character string.
<i>WcString2</i>	Specifies a wide-character string.
<i>Number</i>	Specifies the range of characters to process.

## Related Information

The **mbsncat** subroutine, **mbsncmp** subroutine, **mbsncpy** subroutine, **wcscat** (“wcscat, wcschr, wcsncmp, wcsncpy, or wcsncpy Subroutine” on page 444) subroutine, **wcscmp** (“wcscat, wcschr, wcsncmp, wcsncpy, or wcsncpy Subroutine” on page 444) subroutine, **wcscpy** (“wcscat, wcschr, wcsncmp, wcsncpy, or wcsncpy Subroutine” on page 444) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview, Multibyte and Wide Character String Comparison Subroutines, and Wide Character String Copy Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## wcspbrk Subroutine

### Purpose

Locates the first occurrence of characters in a string.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
wchar_t *wcspbrk( WcString1, WcString2)
const wchar_t *WcString1;
const wchar_t *WcString2;
```

### Description

The **wcspbrk** subroutine locates the first occurrence in the wide character string pointed to by the *WcString1* parameter of any wide character from the string pointed to by the *WcString2* parameter.

### Parameters

*WcString1*            Points to a wide-character string being searched.  
*WcString2*            Points to a wide-character string.

### Return Values

If no **wchar\_t** character from the *WcString2* parameter occurs in the *WcString1* parameter, the **wcspbrk** subroutine returns a pointer to the wide character, or a null value.

### Related Information

The **mbspbrk** subroutine, **wcschr** (“wcscat, wcschr, wcscmp, wcsncpy, or wcsncpy Subroutine” on page 444) subroutine, **wcscspn** (“wcscat, wcschr, wcscmp, wcsncpy, or wcsncpy Subroutine” on page 444) subroutine, **wcsrchr** (“wcsrchr Subroutine”) subroutine, **wcsspn** (“wcsspn Subroutine” on page 453) subroutine, **wcstok** (“wcstok Subroutine” on page 458) subroutine, **wcswcs** (“wcswcs Subroutine” on page 464) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Wide Character String Search Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## wcsrchr Subroutine

### Purpose

Locates a **wchar\_t** character in a wide-character string.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <wctr.h>
```

```
wchar_t *wcsrchr ( WcString, WideCharacter)  
const wchar_t *WcString;  
wint_t WideCharacter;
```

## Description

The **wcsrchr** subroutine locates the last occurrence of the *WideCharacter* value in the string pointed to by the *WcString* parameter. The terminating **wchar\_t** null character is considered to be part of the string.

## Parameters

<i>WcString</i>	Points to a string.
<i>WideCharacter</i>	Specifies a <b>wchar_t</b> character.

## Return Values

The **wcsrchr** subroutine returns a pointer to the *WideCharacter* parameter value, or a null pointer if that value does not occur in the specified string.

## Related Information

The **mbschr** subroutine, **mbsrchr** subroutine, **wcschr** (“wscat, wcschr, wcsncmp, wcsncpy, or wcsncpy Subroutine” on page 444) subroutine, **wcscspn** (“wscat, wcschr, wcsncmp, wcsncpy, or wcsncpy Subroutine” on page 444) subroutine, **wcspbrk** (“wcspbrk Subroutine” on page 451) subroutine, **wcsspn** (“wcsspn Subroutine” on page 453) subroutine, **wcstok** (“wcstok Subroutine” on page 458) subroutine, **wcswcs** (“wcswcs Subroutine” on page 464) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Understanding Wide Character String Search Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## wcsrtombs Subroutine

### Purpose

Convert a wide-character string to a character string (restartable).

### Library

Standard library (**libc.a**)

### Syntax

```
#include <wchar.h>  
size_t wcsrtombs (char * dst, const wchar_t ** src, size_t len, mbstate_t * ps);
```

## Description

The **wcsrtombs** function converts a sequence of wide-characters from the array indirectly pointed to by **src** into a sequence of corresponding characters, beginning in the conversion state described by the object pointed to by **ps**. If **dst** is not a null pointer, the converted characters are then stored into the array pointed to by **dst**. Conversion continues up to and including a terminating null wide-character, which is also stored. Conversion stops earlier in the following cases:

- When a code is reached that does not correspond to a valid character.
- When the next character would exceed the limit of **len** total bytes to be stored in the array pointed to by **dst** (and **dst** is not a null pointer).

Each conversion takes place as if by a call to the **wcrtomb** function.

If **dst** is not a null pointer, the pointer object pointed to by **src** 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.

If **ps** is a null pointer, the **wcsrtombs** function uses its own internal **mbstate\_t** object, which is initialised at program startup to the initial conversion state. Otherwise, the **mbstate\_t** object pointed to by **ps** is used to completely describe the current conversion state of the associated character sequence. The implementation will behave as if no function defined in this specification calls **wcsrtombs**.

The behavior of this function is affected by the LC\_CTYPE category of the current locale.

## Return Values

If conversion stops because a code is reached that does not correspond to a valid character, an encoding error occurs. In this case, the **wcsrtombs** function stores the value of the macro **EILSEQ** in **errno** and returns **(size\_t)-1**; the conversion state is undefined. Otherwise, it returns the number of bytes in the resulting character sequence, not including the terminating null (if any).

## Error Codes

The **wcsrtombs** function may fail if:

- EINVAL**     **ps** points to an object that contains an invalid conversion state.  
**EILSEQ**     A wide-character code does not correspond to a valid character.

## Related Information

The **wctomb** (“wctomb Subroutine” on page 468) subroutine.

---

## wcsspn Subroutine

### Purpose

Returns the number of wide characters in the initial segment of a string.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <wcstr.h>
```

```
size_t wcsspncpy( WcString1, WcString2)
const wchar_t *WcString1, *WcString2;
```

## Description

The **wcsspncpy** subroutine computes the number of **wchar\_t** characters in the initial segment of the string pointed to by the *WcString1* parameter. The *WcString1* parameter consists entirely of **wchar\_t** characters from the string pointed to by the *WcString2* parameter.

## Parameters

*WcString1*            Points to the initial segment of a string.  
*WcString2*            Points to a set of characters string.

## Return Values

The **wcsspncpy** subroutine returns the number of **wchar\_t** characters in the segment.

## Related Information

The **wcschr** (“wscat, wcschr, wscmp, wcscopy, or wcsncpy Subroutine” on page 444) subroutine, **wcsncpy** (“wscat, wcschr, wscmp, wcscopy, or wcsncpy Subroutine” on page 444) subroutine, **wcspbrk** (“wcspbrk Subroutine” on page 451) subroutine, **wcsrchr** (“wcsrchr Subroutine” on page 451) subroutine, **wcstok** (“wcstok Subroutine” on page 458) subroutine, **wcswcs** (“wcswcs Subroutine” on page 464) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Wide Character String Search Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## wcsstr Subroutine

### Purpose

Find a wide-character substring.

### Library

Standard library (**libc.a**)

### Syntax

```
#include <wchar.h>
wchar_t *wcsstr (const wchar_t * ws1, const wchar_t * ws2);
```

### Description

The **wcsstr** function locates the first occurrence in the wide-character string pointed to by **ws1** of the sequence of wide-characters (excluding the terminating null wide-character) in the wide-character string pointed to by **ws2**.

### Return Values

On successful completion, **wcsstr** returns a pointer to the located wide-character string, or a null pointer if the wide-character string is not found.

If **ws2** points to a wide-character string with zero length, the function returns **ws1**.

---

## wcstod, wcstof, or wcstold Subroutine

### Purpose

Converts a wide character string to a double-precision number.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
double wcstod ( nptr, endptr )
const wchar_t *nptr;
wchar_t **endptr;
```

```
float wcstof ( nptr, endptr )
const wchar_t *restrict nptr;
wchar_t **restrict endptr;
```

```
long double wcstold ( nptr, endptr )
const wchar_t *restrict format;
wchar_t **restrict nptr;
```

### Description

The **wcstod**, **wcstof**, and **wcstold** subroutines convert the initial portion of the wide-character string pointed to by *nptr* to **double**, **float** and **long double** representation, respectively. First, they decompose the input wide-character string into three parts:

- An initial, possibly empty, sequence of white-space wide-character codes.
- A subject sequence interpreted as a floating-point constant or representing infinity or NaN.
- 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.

Then they convert the subject sequence to a floating-point number, and return the result.

The expected form of the subject sequence is an optional plus or minus sign, and one of the following:

- A non-empty sequence of decimal digits optionally containing a radix character, and an optional exponent part.
- A 0x or 0X, and a non-empty sequence of hexadecimal digits optionally containing a radix character, and an optional binary exponent part.
- One of INF or INFINITY, or any other wide string equivalent except for case.
- One of NAN or NAN(*n-wchar-sequence* *opt*), or any other wide string ignoring case in the NAN part, where:

```
n-wchar-sequence:
  digit
  nondigit
  n-wchar-sequence digit
  n-wchar-sequence nondigit
```

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) are 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. If neither an exponent part or 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, if representable in the return type, or else as if it were a floating constant that is too large for the range of the return type. A wide-character sequence NAN or NAN(*n-wchar-sequence<sub>opt</sub>*) is interpreted as a quiet NaN, if supported in the return type, or else as if it were a subject sequence part that does not have the expected form. The meaning of the *n-wchar* sequences is implementation-defined. 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 the hexadecimal form and FLT\_RADIX is a power of 2, the conversion will be rounded in an implementation-defined manner.

The radix character is as 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.

In other than the C or POSIX locales, other implementation-defined subject sequences may be accepted.

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**, **wcstof**, and **wcstold** subroutines do not change the setting of the **errno** global variable if successful.

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, call **wcstod**, **wcstof**, or **wcstold**, and check **errno**.

## Parameters

<i>nptr</i>	Contains a pointer to the wide character string to be converted to a double-precision value.
<i>endptr</i>	Contains a pointer to the position in the string specified by the <i>nptr</i> parameter where a wide character is found that is not a valid character for the purpose of this conversion.

## Return Values

Upon successful completion, the **wcstod**, **wcstof**, and **wcstold** subroutines return the converted value. If no conversion could be performed, 0 is returned and the **errno** global variable may be set to **EINVAL**.

If the correct value is outside the range of representable values, plus or minus **HUGE\_VAL**, **HUGE\_VALF**, or **HUGE\_VALL** is returned (according to the sign of the value), and **errno** is set to ERANGE.

If the correct value would cause underflow, a value whose magnitude is no greater than the smallest normalized positive number in the return type is returned and **errno** set to ERANGE.

## Related Information

“scanf, fscanf, sscanf, or wsscanf Subroutine” on page 109, “setlocale Subroutine” on page 136, and “strtol, strtoul, strtoll, strtoull, or atoi Subroutine” on page 299.

cctype, isalpha, isupper, islower, isdigit, isxdigit, isalnum, isspace, ispunct, isprint, isgraph, iscntrl, or isascii Subroutines and localeconv Subroutine in *AIX 5L Version 5.2 Technical Reference: Base Operating System and Extensions Volume 1*.

---

## wcstoimax or wcstoumax Subroutine

### Purpose

Converts a wide-character string to an integer type.

### Syntax

```
#include <stddef.h>
#include <inttypes.h>

intmax_t wcstoimax (nptr, endptr, base)
const wchar_t *restrict nptr;
wchar_t **restrict endptr;
int base;

uintmax_t wcstoumax (nptr, endptr, base)
const wchar_t *restrict nptr;
wchar_t **restrict endptr;
int base;
```

### Description

The **wcstoimax** or **wcstoumax** subroutines are equivalent to the **wcstol**, **wcstoll**, **wcstoul**, and **wcstoull** subroutines, respectively, except that the initial portion of the wide string is converted to **intmax\_t** and **uintmax\_t** representation, respectively.

### Parameters

<i>nptr</i>	Points to the wide-character string.
<i>endptr</i>	Points to the object where the final wide-character string is stored.
<i>base</i>	Determines the subject sequence interpreted as an integer.

### Return Values

The **wcstoimax** or **wcstoumax** subroutines return the converted value, if any.

If no conversion could be performed, zero is returned. If the correct value is outside the range of representable values, **{INTMAX\_MAX}**, **{INTMAX\_MIN}**, or **{UINTMAX\_MAX}** is returned (according to the return type and sign of the value, if any), and the **errno** global variable is set to ERANGE.

### Related Information

The “wcstol or wcstoll Subroutine” on page 459.

inttypes.h in *AIX 5L Version 5.2 Files Reference*.

---

## wcstok Subroutine

### Purpose

Converts wide-character strings to tokens.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <wchar.h>
```

```
wchar_t *wcstok ( WcString1, WcString2, ptr)
wchar_t *WcString1;
const wchar_t *WcString2;
wchar_t **ptr
```

### Description

A sequence of calls to the **wcstok** subroutine breaks the wide-character string pointed to by *WcString1* into a sequence of tokens, each of which is delimited by a wide-character code from the wide-character string pointed to by *WcString2*. The third argument points to a caller-provided **wchar\_t** pointer where **wcstok** stores information necessary for it to continue scanning the same wide-character string.

The first call in the sequence has *WcString1* as its first argument and is followed by calls with a nullpointer as their first argument. The separator string pointed to by *WcString2* may be different from call to call.

The first call in the sequence searches the wide-character string pointed to by *WcString1* for the first wide-character code that is not contained in the current separator string pointed to by *WcString2*. If no such wide-character code is found, then there are no tokens in the wide-character string pointed to by *WcString1* and **wcstok** returns a null pointer. If such a wide-character code is found, it is the start of the first token.

The **wcstok** subroutine then searches from there for a wide-character code that is 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 *WcString1*, and subsequent searches for a token returns 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 **wcstok** subroutine saves a pointer to the following wide-character code, from which the next search for a token starts.

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.

The implementation behaves as if no function calls **wcstok**.

### Parameters

<i>ptr</i>	Contains a pointer to a caller-provided <b>wchar_t</b> pointer where <b>wcstok</b> stores information necessary for it to continue scanning the same wide-character string.
<i>WcString1</i>	Contains a pointer to the wide-character string to be searched.
<i>WcString2</i>	Contains a pointer to the string of wide-character token delimiters.

## Return Values

Upon successful completion, **wcstok** returns a pointer to the first wide-character code of a token. Otherwise, if there is no token, **wcstok** returns a null pointer.

## Examples

To convert a wide-character string to tokens, use the following:

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

main()
{
    wchar_t *WCString1 = L"?a???b,,#c";
    wchar_t *ptr;
    wchar_t *pwcs;

    (void)setlocale(LC_ALL, "");
    pwcs = wcstok(WCString1, L"?", &ptr);
        /* pwcs points to the token L"a"*/
    pwcs = wcstok((wchar_t *)NULL, L",", &ptr);
        /* pwcs points to the token L"??b"*/
    pwcs = wcstok( (wchar_t *)NULL, L"#,", &ptr);
        /* pwcs points to the token L"c"*/
}
```

## Related Information

The **wcschr** (“wscat, wcschr, wcsncmp, wcsncpy, or wcsncpy Subroutine” on page 444) subroutine, **wcscspn** (“wscat, wcschr, wcsncmp, wcsncpy, or wcsncpy Subroutine” on page 444) subroutine, **wcspbrk** (“wcspbrk Subroutine” on page 451) subroutine, **wcsrchr** (“wcsrchr Subroutine” on page 451) subroutine, **wcsspn** (“wcsspn Subroutine” on page 453) subroutine, **wcstod** (“wcstod, wcstof, or wcstold Subroutine” on page 455) subroutine, **wcstol** (“wcstol or wcstoll Subroutine”) subroutine, **wcstoul** (“wcstoul or wcstoull Subroutine” on page 463) subroutine, **wcswcs** (“wcswcs Subroutine” on page 464) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview for Programming and Wide Character String Search Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## wcstol or wcstoll Subroutine

### Purpose

Converts a wide-character string to a long integer representation.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
long int wcstol ( Nptr, Endptr, Base)
const wchar_t *Nptr;
wchar_t **Endptr;
int Base;
```

```
long long int wcstoll (*Nptr, **Endptr, Base)
const wchar_t *Nptr;
wchar_t **Endptr;
int Base
```

## Description

The **wcstol** subroutine converts a wide-character string to a long integer representation. The **wcstoll** subroutine converts a wide-character string to a long long integer representation.

1. An initial, possibly empty, sequence of white-space wide-character codes (as specified by the **iswspace** subroutine)
2. A subject sequence interpreted as an integer and represented in a radix determined by the *Base* parameter
3. A final wide-character string of one or more unrecognized wide-character codes, including the terminating wide-character null of the input wide-character string

If possible, the subject is then converted to an integer, and the result is returned.

The *Base* parameter can take the following values: 0 through 9, or a (or A) through z (or Z). There are potentially 36 values for the base. If the base value is 0, the expected form of the subject string is that of a decimal, octal, or hexadecimal constant, any of which can be preceded by a + (plus sign) or - (minus sign). A decimal constant starts with a non zero digit, and is composed 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. A hexadecimal constant is defined as the prefix 0x (or 0X) followed by a sequence of decimal digits and the letters a (or A) to f (or F) with values ranging from 10 (for a or A) to 15 (for f or F).

If the base value is between 2 and 36, the expected form of the subject sequence is a sequence of letters and digits representing an integer in the radix specified by the *Base* parameter, optionally preceded by a + or -, but not including an integer suffix. The letters a (or A) through z (or Z) are ascribed the values of 10 to 35. Only letters whose values are less than that of the base are permitted. If the value of base is 16, the characters 0x or 0X may optionally precede the sequence of letters or digits, following the sign, if present.

The wide-character string is parsed to skip the initial space characters (as determined by the **iswspace** subroutine). Any non-space character signifies the start of a subject string that may form an integer in the radix specified by the *Base* parameter. The subject sequence is defined to be the longest initial substring that is a long integer of the expected form. Any character not satisfying this form begins the final portion of the wide-character string pointed to by the *Endptr* parameter on return from the call to the **wcstol** or **wcstoll** subroutine.

## Parameters

<i>Nptr</i>	Contains a pointer to the wide-character string to be converted to a long integer number.
<i>Endptr</i>	Contains a pointer to the position in the <i>Nptr</i> parameter string where a wide-character is found that is not a valid character.
<i>Base</i>	Specifies the radix in which the characters are interpreted.

## Return Values

The **wcstol** and **wcstoll** subroutines return the converted value of the long or long long integer if the expected form is found. If no conversion could be performed, a value of 0 is returned. If the converted value is outside the range of representable values, **LONG\_MAX** or **LONG\_MIN** is returned for the **wcstol** subroutine and **LLONG\_MAX** or **LLONG\_MIN** is returned for the **wcstoll** subroutine (according to the sign of the value). The value of **errno** is set to **ERANGE**. If the base value specified by the *Base* parameter is not supported, **EINVAL** is returned.

If the subject sequence has the expected form, it is interpreted as an integer constant in the appropriate base. A pointer to the final string is stored in the *Endptr* parameter if that parameter is not a null pointer.

If the subject sequence is empty or does not have a valid form, no conversion is done. The value of the *Nptr* parameter is stored in the *Endptr* parameter if that parameter is not a null pointer.

Since 0, **LONG\_MIN**, and **LONG\_MAX** (for **wcstol**) and **LLONG\_MIN**, and **LLONG\_MAX** (for **wcstoll**) are returned in the event of an error and are also valid returns if the **wcstol** or **wcstoll** subroutine is successful, applications should set the **errno** global variable to 0 before calling either subroutine, and check **errno** after return. If the **errno** global value has changed, an error occurred.

## Examples

To convert a wide-character string to a signed long integer, use the following code:

```
#include <stdlib.h>
#include <locale.h>
#include <errno.h>
main()
{
    wchar_t *WCString, *endptr;
    long int retval;
    (void)setlocale(LC_ALL, "");
    /**Set errno to 0 so a failure for wcstol can be
    **detected */
    errno=0;
    /*
    **Let WCString point to a wide character null terminated
    ** string containing a signed long integer value
    **
    */
    /*retval = wcstol ( WCString &endptr, 0 );
    /* Check errno, if it is non-zero, wcstol failed */
    if (errno != 0) {
        /*Error handling*/
    }
    else if (&WCString == endptr) {
        /* No conversion could be performed */
        /* Handle this case accordingly. */
    }
    /* retval contains long integer */
}
```

## Related Information

The **iswspace** subroutine, **wcstod** (“wcstod, wcstof, or wcstold Subroutine” on page 455) subroutine, **wcstoul** (“wcstoul or wcstoull Subroutine” on page 463) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Wide Character String Conversion Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## wcstombs Subroutine

### Purpose

Converts a sequence of wide characters into a sequence of multibyte characters.

## Library

Standard C Library (**libc.a**)

## Syntax

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

```
size_t wcstombs ( String, WcString, Number)  
char *String;  
const wchar_t *WcString;  
size_t Number;
```

## Description

The **wcstombs** subroutine converts the sequence of wide characters pointed to by the *WcString* parameter to a sequence of corresponding multibyte characters and places the results in the area pointed to by the *String* parameter. The conversion is terminated when the null wide character is encountered or when the number of bytes specified by the *Number* parameter (or the value of the *Number* parameter minus 1) has been placed in the area pointed to by the *String* parameter. If the amount of space available in the area pointed to by the *String* parameter would cause a partial multibyte character to be stored, the subroutine uses a number of bytes equalling the value of the *Number* parameter minus 1, because only complete multibyte characters are allowed.

## Parameters

<i>String</i>	Points to the area where the result of the conversion is stored. If the <i>String</i> parameter is a null pointer, the subroutine returns the number of bytes required to hold the conversion.
<i>WcString</i>	Points to a wide-character string.
<i>Number</i>	Specifies a number of bytes to be converted.

## Return Values

The **wcstombs** subroutine returns the number of bytes modified. If a wide character is encountered that is not valid, a value of -1 is returned.

## Error Codes

The **wcstombs** subroutine is unsuccessful if the following error occurs:

<b>EILSEQ</b>	An invalid character sequence is detected, or a wide-character code does not correspond to a valid character.
---------------	---

## Related Information

The **mbstowcs** subroutine, **mbtowc** subroutine, **wcslen** (“wcslen Subroutine” on page 449) subroutine, **wctomb** (“wctomb Subroutine” on page 468) subroutine.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview and Multibyte Code and Wide Character Code Conversion Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

---

## wcstoul or wcstoull Subroutine

### Purpose

Converts wide character strings to unsigned long or long long integer representation.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
unsigned long int wcstoul (Nptr, Endptr, Base)
```

```
const wchar_t * Nptr;
```

```
wchar_t ** Endptr;
```

```
int Base;
```

```
unsigned long long int wcstoull (Nptr, Endptr, Base)
```

```
const wchar_t *Nptr;
```

```
wchar_t **Endptr;
```

```
int Base;
```

### Description

The **wcstoul** and **wcstoull** subroutines convert the initial portion of the wide character string pointed to by the *Nptr* parameter to an unsigned long or long long integer representation. To do this, it parses the wide character string pointed to by the *Nptr* parameter to obtain a valid string (that is, subject string) for the purpose of conversion to an unsigned long integer. It then points the *Endptr* parameter to the position where an unrecognized character, including the terminating null, is found.

The base specified by the *Base* parameter can take the following values: 0 through 9, a (or A) through z (or Z). There are potentially 36 values for the base. If the base value is 0, the expected form of the subject string is that of an unsigned integer constant, with an optional + (plus sign) or - (minus sign), but not including the integer suffix. If the base value 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 the *Base* parameter, optionally preceded by a + or -, but not including an integer suffix.

The letters a (or A) through z (or Z) are ascribed the values of 10 to 35. Only letters whose values are less than that of the base are permitted. If the value of the base is 16, the characters 0x (or 0X) may optionally precede the sequence of letters or digits, following a + or - . present.

The wide character string is parsed to skip the initial white-space characters (as determined by the **iswspace** subroutine). Any nonspace character signifies the start of a subject string that may form an unsigned long integer in the radix specified by the *Base* parameter. The subject sequence is defined to be the longest initial substring that is an unsigned long integer of the expected form. Any character not satisfying this expected form begins the final portion of the wide character string pointed to by the *Endptr* parameter on return from the call to this subroutine.

### Parameters

<i>Nptr</i>	Contains a pointer to the wide character string to be converted to an unsigned long integer.
<i>Endptr</i>	Contains a pointer to the position in the <i>Nptr</i> string where a wide character is found that is not a valid character for the purpose of this conversion.
<i>Base</i>	Specifies the radix in which the wide characters are interpreted.

## Return Values

The **wcstoul** and **wcstoull** subroutines return the converted value of the unsigned long or long long integer if the expected form is found. If no conversion could be performed, a value of 0 is returned. If the converted value is outside the range of representable values, a **ULONG\_MAX** value is returned (for **wcstoul**), and **ULLONG\_MAX** is returned (for **wcstoull**), and the value of the **errno** global variable is set to a **ERANGE** value.

If the subject sequence has the expected form, it is interpreted as an integer constant in the appropriate base. A pointer to the final string is stored in the *Endptr* parameter if that parameter is not a null pointer. If the subject sequence is empty or does not have a valid form, no conversion is done and the value of the *Nptr* parameter is stored in the *Endptr* parameter if it is not a null pointer.

If the radix specified by the *Base* parameter is not supported, an **EINVAL** value is returned. If the value to be returned is not representable, an **ERANGE** value is returned.

## Examples

To convert a wide character string to an unsigned long integer, use the following code:

```
#include <stdlib.h>
#include <locale.h>
#include <errno.h>
extern int errno;
main()
{
    wchar_t *WCString, *EndPtr;
    unsigned long int  retval;
    (void)setlocale(LC_ALL, "");
    /*
    ** Let WCString point to a wide character null terminated
    ** string containing an unsigned long integer value.
    **
    */
    retval = wcstoul ( WCString &EndPtr, 0 );
    if(retval==0) {
        /* No conversion could be performed */
        /* Handle this case accordingly. */
    } else if(retval == ULONG_MAX) {
        /* Error handling */
    }
    /* retval contains the unsigned long integer value. */
}
```

## Related Information

National Language Support Overview and Wide Character String Conversion Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*

---

## wcswcs Subroutine

### Purpose

Locates first occurrence of a wide character in a string.

### Library

Standard C Library (**libc.a**)

## Syntax

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

```
wchar_t *wswcs( WcString1, WcString2)  
const wchar_t *WcString1, *WcString2;
```

## Description

The **wswcs** subroutine locates the first occurrence, in the string pointed to by the *WcString1* parameter, of a sequence of **wchar\_t** characters (excluding the terminating **wchar\_t** null character) from the string pointed to by the *WcString2* parameter.

## Parameters

<i>WcString1</i>	Points to the wide-character string being searched.
<i>WcString2</i>	Points to a wide-character string, which is a source string.

## Return Values

The **wswcs** subroutine returns a pointer to the located string, or a null value if the string is not found. If the *WcString2* parameter points to a string with 0 length, the function returns the *WcString1* value.

## Related Information

The **mbspbrk** subroutine, **wcschr** (“**wcscat**, **wcschr**, **wscmp**, **wscopy**, or **wcscspn** Subroutine” on page 444) subroutine, **wcscspn** (“**wcscat**, **wcschr**, **wscmp**, **wscopy**, or **wcscspn** Subroutine” on page 444) subroutine, **wcspbrk** (“**wcspbrk** Subroutine” on page 451) subroutine, **wcsrchr** (“**wcsrchr** Subroutine” on page 451) subroutine, **wcsspn** (“**wcsspn** Subroutine” on page 453) subroutine, **wcstok** (“**wcstok** Subroutine” on page 458) subroutine.

National Language Support Overview and Wide Character String Search Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*

---

## wcswidth Subroutine

### Purpose

Determines the display width of wide character strings.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int wcswidth (* PwCS, n)  
const wchar_t *PwCS;  
size_t n;
```

## Description

The **wcswidth** subroutine determines the number of display columns to be occupied by the number of wide characters specified by the *N* parameter in the string pointed to by the *Pwcs* parameter. The **LC\_CTYPE** category affects the behavior of the **wcswidth** subroutine. Fewer than the number of wide characters specified by the *N* parameter are counted if a null character is encountered first.

## Parameters

*N* Specifies the maximum number of wide characters whose display width is to be determined.  
*Pwcs* Contains a pointer to the wide character string.

## Return Values

The **wcswidth** subroutine returns the number of display columns to be occupied by the number of wide characters (up to the terminating wide character null) specified by the *N* parameter (or fewer) in the string pointed to by the *Pwcs* parameter. A value of zero is returned if the *Pwcs* parameter is a wide character null pointer or a pointer to a wide character null (that is, *Pwcs* or *\*Pwcs* is null). If the *Pwcs* parameter points to an unusable wide character code, -1 is returned.

## Examples

To find the display column width of a wide character string, use the following:

```
#include <string.h>
#include <locale.h>
#include <stdlib.h>
main()
{
    wchar_t *pwcs;
    int      retval, n ;
    (void)setlocale(LC_ALL, "");
    /* Let pwcs point to a wide character null terminated
    ** string. Let n be the number of wide characters whose
    ** display column width is to be determined.
    */
    retval= wcswidth( pwcs, n );
    if(retval == -1){
        /* Error handling. Invalid wide character code
        ** encountered in the wide character string pwcs.
        */
    }
}
```

## Related Information

The **wcwidth** (“wcwidth Subroutine” on page 471) subroutine.

National Language Support Overview and Wide Character Display Column Width Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## wcsxfrm Subroutine

### Purpose

Transforms wide-character strings to wide-character codes of current locale.

## Library

Standard C Library (**libc.a**)

## Syntax

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

```
size_t wcsxfrm ( WcString1, WcString2, Number)  
wchar_t *WcString1;  
const wchar_t *WcString2;  
size_t Number;
```

## Description

The **wcsxfrm** subroutine transforms the wide-character string specified by the *WcString2* parameter into a string of wide-character codes, based on the collation values of the wide characters in the current locale as specified by the **LC\_COLLATE** category. No more than the number of character codes specified by the *Number* parameter are copied into the array specified by the *WcString1* parameter. When two such transformed wide-character strings are compared using the **wscmp** subroutine, the result is the same as that obtained by a direct call to the **wscoll** subroutine on the two original wide-character strings.

## Parameters

<i>WcString1</i>	Points to the destination wide-character string.
<i>WcString2</i>	Points to the source wide-character string.
<i>Number</i>	Specifies the maximum number of wide-character codes to place into the array specified by <i>WcString1</i> . To determine the necessary size specification, set the <i>Number</i> parameter to a value of 0, so that the <i>WcString1</i> parameter becomes a null pointer. The return value plus 1 is the size necessary for the conversion.

## Return Values

If the *WcString1* parameter is a wide-character null pointer, the **wcsxfrm** subroutine returns the number of wide-character elements (not including the wide-character null terminator) required to store the transformed wide character string. If the count specified by the *Number* parameter is sufficient to hold the transformed string in the *WcString1* parameter, including the wide character null terminator, the return value is set to the actual number of wide character elements placed in the *WcString1* parameter, not including the wide character null. If the return value is equal to or greater than the value specified by the *Number* parameter, the contents of the array pointed to by the *WcString1* parameter are indeterminate. This occurs whenever the *Number* value parameter is too small to hold the entire transformed string. If an error occurs, the **wcsxfrm** subroutine returns the **size\_t** data type with a value of -1 and sets the **errno** global variable to indicate the error.

If the wide character string pointed to by the *WcString2* parameter contains wide character codes outside the domain of the collating sequence defined by the current locale, the **wcsxfrm** subroutine returns a value of **EINVAL**.

## Related Information

The **wscmp** (“wscat, wcschr, wscmp, wcsncpy, or wcsncpy Subroutine” on page 444) subroutine, **wscoll** (“wscoll Subroutine” on page 446) subroutine.

National Language Support Overview and Wide Character String Collation Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## wctob Subroutine

### Purpose

Wide-character to single-byte conversion.

### Library

Standard library (**libc.a**)

### Syntax

```
#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.

### 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.

### Related Information

The **btowc** subroutine.

---

## wctomb Subroutine

### Purpose

Converts a wide character into a multibyte character.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>

int wctomb (Storage, WideCharacter)
char *Storage;
wchar_t WideCharacter;
```

### Description

The **wctomb** subroutine determines the number of bytes required to represent the wide character specified by the *WideCharacter* parameter as the corresponding multibyte character. It then converts the *WideCharacter* value to a multibyte character and stores the results in the area pointed to by the *Storage* parameter. The **wctomb** subroutine can store a maximum of **MB\_CUR\_MAX** bytes in the area pointed to by the *Storage* parameter. Thus, the length of the area pointed to by the *Storage* parameter should be at least **MB\_CUR\_MAX** bytes. The **MB\_CUR\_MAX** macro is defined in the **stdlib.h** file.

## Parameters

*Storage* Points to an area where the result of the conversion is stored.  
*WideCharacter* Specifies a wide-character value.

## Return Values

The **wctomb** subroutine returns a 0 if the *Storage* parameter is a null pointer. If the *WideCharacter* parameter does not correspond to a valid multibyte character, a -1 is returned. Otherwise, the number of bytes that comprise the multibyte character is returned.

## Related Information

The **mbtowc** subroutine, **mbstowcs** subroutine, **wcslen** (“wcslen Subroutine” on page 449) subroutine, **wcstombs** (“wcstombs Subroutine” on page 461) subroutine.

National Language Support Overview and Multibyte Code and Wide Character Code Conversion Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## wctrans Subroutine

### Purpose

Define character mapping.

### Library

Standard library (**libc.a**)

### Syntax

```
#include <wctype.h>
wctrans_t wctrans (const char * charclass);
```

### Description

The **wctrans** function is defined for valid character mapping names identified in the current locale. The **charclass** 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”.

The function returns a value of type **wctrans\_t**, which can be used as the second argument to subsequent calls of **towctrans**. The **wctrans** function determines values of **wctrans\_t** according to the rules of the coded character set defined by character mapping information in the program’s locale (category LC\_CTYPE). The values returned by **wctrans** are valid until a call to **setlocale** that modifies the category LC\_CTYPE.

### Return Values

The **wctrans** function returns 0 if the given character mapping name is not valid for the current locale (category LC\_CTYPE), otherwise it returns a non-zero object of type **wctrans\_t** that can be used in calls to **towctrans**.

## Error Codes

The **wctrans** function may fail if:

**EINVAL** The character mapping name pointed to by `charclass` is not valid in the current locale.

## Related Information

The **towctrans** (“`towctrans` Subroutine” on page 365) subroutine.

---

## wctype or get\_wctype Subroutine

### Purpose

Obtains a handle for valid property names in the current locale for wide characters.

### Library

Standard C library (**libc.a**).

### Syntax

```
#include <wchar.h>
```

```
wctype_t wctype ( Property)  
const char *Property;
```

```
wctype_t get_wctype ( Property)  
char *Property;
```

### Description

The **wctype** subroutine obtains a handle for valid property names for wide characters as defined in the current locale. The handle is of data type **wctype\_t** and can be used as the *WC\_PROP* parameter in the **iswctype** subroutine. Values returned by the **wctype** subroutine are valid until the **setlocale** subroutine modifies the **LC\_CTYPE** category. The **get\_wctype** subroutine is identical to the **wctype** subroutine.

The **wctype** subroutine adheres to X/Open Portability Guide Issue 5.

## Parameters

<i>Property</i>	Points to a string that identifies a generic character class for which code set-specific information is required. The basic character classes are:
<b>alnum</b>	Alphanumeric character.
<b>alpha</b>	Alphabetic character.
<b>blank</b>	Space and tab characters.
<b>cntrl</b>	Control character. No characters in <b>alpha</b> or <b>print</b> are included.
<b>digit</b>	Numeric digit character.
<b>graph</b>	Graphic character for printing. Does not include the space character or <b>cntrl</b> characters, but does include all characters in <b>digit</b> and <b>punct</b> .
<b>lower</b>	Lowercase character. No characters in <b>cntrl</b> , <b>digit</b> , <b>punct</b> , or <b>space</b> are included.
<b>print</b>	Print character. Includes characters in <b>graph</b> , but does not include characters in <b>cntrl</b> .
<b>punct</b>	Punctuation character. No characters in <b>alpha</b> , <b>digit</b> , or <b>cntrl</b> , or the space character are included.
<b>space</b>	Space characters.
<b>upper</b>	Uppercase character.
<b>xdigit</b>	Hexadecimal character.

## Return Values

Upon successful completion, the subroutine returns a value of type **wctype\_t**, which is a handle for valid property names in the current locale. Otherwise, it returns a value or - 1 if the *Property* parameter specifies a character class that is not valid for the current locale.

## Related Information

The **iswalnum** subroutine, **iswalpha** subroutine, **iswcntrl** subroutine, **iswctype** subroutine, **iswdigit** subroutine, **iswgraph** subroutine, **iswlower** subroutine, **iswprint** subroutine, **iswpunct** subroutine, **iswspace** subroutine, **iswupper** subroutine, **iswxdigit** subroutine, **setlocale** (“setlocale Subroutine” on page 136) subroutine, **towlower** (“towlower Subroutine” on page 366) subroutine, **towupper** (“towupper Subroutine” on page 367) subroutine.

National Language Support Overview, Wide Character Classification Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## wcwidth Subroutine

### Purpose

Determines the display width of wide characters.

### Library

Standard C Library (**libc.a**)

## Syntax

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

```
int wwidth ( WC)
```

```
wchar_t WC;
```

## Description

The **wwidth** subroutine determines the number of display columns to be occupied by the wide character specified by the *WC* parameter. The **LC\_CTYPE** subroutine affects the behavior of the **wwidth** subroutine.

## Parameters

*WC* Specifies a wide character.

## Return Values

The **wwidth** subroutine returns the number of display columns to be occupied by the *WC* parameter. If the *WC* parameter is a wide character null, a value of 0 is returned. If the *WC* parameter points to an unusable wide character code, -1 is returned.

## Examples

To find the display column width of a wide character, use the following:

```
#include <string.h>
#include <locale.h>
#include <stdlib.h>
main()
{
    wchar_t wc;
    int  retval;
    (void)setlocale(LC_ALL, "");
    /* Let wc be the wide character whose
    ** display width is to be found.
    */
    retval= wwidth( wc );
    if(retval == -1){
        /*
        ** Error handling. Invalid wide character in wc.
        */
    }
}
```

## Related Information

The **wcswidth** (“wcswidth Subroutine” on page 465) subroutine.

National Language Support Overview, Wide Character Display Column Width Subroutines in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## wlm\_assign Subroutine

### Purpose

Manually assigns processes to a class or cancels prior manual assignments for processes.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

```
int wlm_assign ( args)
```

```
struct wlm_assign *args;
```

### Description

The **wlm\_assign** subroutine:

- Assigns a set of processes specified by their process IDs (PIDS) or process group IDs (PGID) to a specified superclass or subclass, thus overriding the automatic class assignment or a prior manual assignment.
- Cancels a previous manual assignment for the specified processes, allowing the processes to be subjected to the automatic assignment rules again.

The target processes are identified by their process ID (pid) or by their process group ID (pgid). The **wlm\_assign** subroutine allows specifying processes using a list of pids, a list of pgids, or both.

The name of a valid superclass or subclass must be specified to manually assign the target processes to a class. If the target class is a superclass, each process is assigned to one of the subclasses of the specified superclass according to the assignment rules for the subclasses of this superclass.

A manual assignment remains in effect (and a process remains in its manually assigned class) until:

- The process terminates.
- The Workload Manager (WLM) is stopped. When WLM is restarted, the manual assignments in effect when WLM was stopped are lost.
- The class the process has been assigned to is deleted.
- The manual assignment for the process is canceled.
- A new manual assignment overrides a prior one.

The name of a valid superclass or subclass must be specified to manually assign the target processes to a class. The assignment can be done or canceled at the superclass level, the subclass level, or both. The interactions between automatic assignment, inheritance and manual assignment are detailed in the Manual Assignment in WLM in *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

Flags in the **wa\_versflags** field described below are used to specify if the requested operation is an assignment or cancellation and at which level.

To assign a process to a class or cancel a prior manual assignment, the caller must have authority both on the process and on the target class. These constraints translate into the following:

- The root user can assign any process to any class.

- A user with administration privileges on the subclasses of a given superclass (that is, the user or group name matches the user or group names specified in the attributes **adminuser** and **admingroup** of the superclass) can manually reassign any process from one of the subclasses of this superclass to another subclass of the superclass.
- A user can manually assign the user's own processes (same real or effective user ID) to a superclass or a subclass, for which the user has manual assignment privileges (that is, the user or group name matches the user or group names specified in the attributes **authuser** and **authgroup** of the superclass or the subclass).

This defines three levels of privilege among the persons who can manually assign processes to classes, root being the highest. For a user to modify or terminate a manual assignment, the user must be at the same level of privilege as the person who issued the last manual assignment, or higher.

**Note:** The **wlm\_assign** subroutine works with the in-core WLM data structures. Even if the WLM current configuration is a set, it applies to the currently loaded regular configuration. If an assignment is made to a class that does not exist in all configurations of the set, it will be lost when the first configuration that does not contain this class is activated (when the class is deleted).

## Parameter

*args* Specifies the address of the **struct wlm\_assign** data structure containing the parameters for the desired class assignment.

The following fields of the **wlm\_args** structure and the embedded substructures can be provided:

<b>wa_versflags</b>	Needs to be initialized with <b>WLM_VERSION</b> . The flags values available, defined in the <b>sys/wlm.h</b> header file, are: <ul style="list-style-type: none"> <li>• <b>WLM_ASSIGN_SUPER</b></li> <li>• <b>WLM_ASSIGN_SUB</b></li> <li>• <b>WLM_ASSIGN_BOTH</b></li> <li>• <b>WLM_UNASSIGN_SUPER</b></li> <li>• <b>WLM_UNASSIGN_SUB</b></li> <li>• <b>WLM_UNASSIGN_BOTH</b></li> </ul>
<b>wa_pids</b>	Specifies the address of the array containing the process IDs of processes to be manually assigned. When this list is empty, a NULL pointer can be passed together with a count of zero (0).
<b>wa_pid_count</b>	Specifies the number of PIDS in the above array. Could be zero (0) if using only pgids to identify the processes.
<b>wa_pgids</b>	Specifies the address of the array containing the process group identifiers (pids) of processes to be manually assigned. When this list is empty, a NULL pointer can be passed together with a count of zero (0).
<b>wa_pgid_count</b>	Specifies the number of PGIDs in the above array. Could be zero (0) if using only pids to identify the processes. If both pids and pgids counts are zero (0), no process is assigned, but the operation is considered successful.
<b>wa_classname</b>	Specifies the full name of the superclass ( <b>super_name</b> ) or the subclass ( <b>super_name.sub_name</b> ) of the class you want to manually assign processes to. The class name field is ignored when canceling an existing manual assignment.

## Return Values

Upon successful completion, the **wlm\_assign** subroutine returns a value of 0. If the **wlm\_assign** subroutine is unsuccessful, a non-0 value is returned. The routine is considered successful if some of the target processes are not found, (to account for process terminations) or are not assigned/deassigned due to a lack of privileges, for instance. If none of the processes in the lists can be assigned/deassigned, this is considered an error.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the **wlm.h** header file.

## Related Information

Manual Assignment in WLM and Workload Manager Application Programming Interface in *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

---

## wlm\_change\_class Subroutine

### Purpose

Changes some of the attributes of a class.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

```
int wlm_change_class ( wlmargs )
```

```
struct wlm_args *wlmargs;
```

### Description

The **wlm\_change\_class** subroutine changes attributes of an existing superclass or subclass. Except for its name, any of the attributes of the class can be modified by a call to **wlm\_change\_class**.

- If the name of a valid configuration is passed in the **confdir** field, the subroutine updates the Workload Manager (WLM) properties files for the target configuration.
- If a null string ('\0') is passed in the **confdir** field, the changes are applied only to the in-core WLM data. No WLM properties file is updated.

The structure of type **struct class\_definition**, which is part of **struct wlm\_args**, has normally been initialized with a call to **wlm\_init\_class\_definition**. Once this has been done, initialize the required fields of this structure (such as the name of the class to be modified) and the fields corresponding to the class attributes you want to modify. For a description of the possible values for the various class attributes and their default values, refer to the description of **wlm.h** in the *AIX 5L Version 5.2 Files Reference*.

The caller must have root authority to change the attributes of a superclass and must have administrator authority on a superclass to change the attributes of a subclass of the superclass.

**Note:** Do not specify a set in the *confdir* field of the **wlm\_args** structure. The **wlm\_change\_class** subroutine cannot apply to a set of time-based configurations.

## Parameters

**wlmargs** Specifies the address of the **struct wlm\_args** data structure containing the **class\_definition** structure for the class to be modified.

The following fields of the **wlm\_args** structure and the embedded substructures need to be provided:

**versflags** Needs to be initialized with **WLM\_VERSION**.

**confdir** Specifies the name of the WLM configuration the target class belongs to. It must be either the name of a valid subdirectory of **/etc/wlm** or an empty string (starting with **\0**).

If the name is a valid subdirectory, the relevant class description file in the given configuration are modified.

If the name is a null string, no description files are updated. The modified class attributes are passed to the kernel similarly to a call to **wlm\_load**.

**name** Specifies the name of the superclass or of the subclass to be modified. If this is a subclass name, it must be of the form **super\_name.sub\_name**. There is no default for this field.

All the other fields can be left at their initial value as set by **wlm\_init\_class\_definition** if the user does not wish to change the current values.

## Return Values

Upon successful completion, the **wlm\_change\_class** subroutine returns a value of 0. If the **wlm\_change\_class** subroutine is unsuccessful, a nonzero value is returned.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the **wlm.h** header file.

## Related Information

The **wlm.h** header file.

The **wlm\_create\_class** (“**wlm\_create\_class** Subroutine” on page 480) subroutine, **wlm\_delete\_class** (“**wlm\_delete\_class** Subroutine” on page 482) subroutine.

Workload Manager Application Programming Interface in *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

---

## wlm\_check subroutine

### Purpose

Check a WLM configuration.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

```
int wlm_check ( config)
```

```
char *config;
```

## Description

The **wlm\_check** subroutine checks the class definitions and the coherency of the assignment rules file(s) (syntax, existence of the classes, validity of user and group names, application path names, etc.) for the configuration whose name is passed as an argument.

If *config* is a null pointer or points to an empty string, **wlm\_check** performs the checks on the configuration files, in the configuration pointed to by **/etc/wlm/current**.

The **wlm\_check** subroutine can apply to a configuration set. If *config* is a configuration set name (or if *config* is not provided and *current* is a configuration set), the checks mentioned above are performed on all configurations of the set, after checking the set itself.

## Parameter

*config*

A pointer to a character string. This pointer should be:

- The address of a character string representing the name of a valid configuration (a subdirectory of **/etc/wlm**)
- A null pointer
- A pointer to a null string ("")

If *config* is a null pointer or a pointer to a null string, the configuration files in the directory pointed to by **/etc/wlm/current** (active configuration) is checked for errors. Otherwise, the configuration files in directory **/etc/wlm/<config\_name>** is checked.

## Return Values

Upon successful completion, a value of 0 is returned. If the **wlm\_checks** subroutine is unsuccessful a non 0 value is returned.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the header file **sys/wlm.h**.

## Related Information

The wlm.h header file.

System Management Concepts: Operating System and Devices, Chapter 13 Workload Manager, **Automatic class Assignment**.

The rules file.

---

## wlm\_classify Subroutine

### Purpose

Determines which classes a process is assigned to.

### Library

Workload Manager Library (**libwlm.a**)

## Syntax

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

int wlm_classify ( config, attributes, class, len)

char *config;

char *attributes;

char *class;

int *len;
```

## Description

The **wlm\_classify** subroutine must be passed the name of a valid configuration and a set of process *attributes* in a format identical to the format of the **rules** file (assignment rules). The names of the classes are copied into the area pointed to by *class*. The integer pointed to by *len* contains the size of the *class* names area on input and the number of matches on output. If the area pointed to by *class* is not big enough to contain the names of all the potential matches, an error is returned.

The normal use of the **wlm\_classify** routine is to explicitly provide all the process classification attributes: **user name**, **group name**, **application pathname**, **type**, and **tag** when applicable. This gives a match to a single class. To implement "what if" scenarios, the interface allows you to leave some of the attributes unspecified by using a hyphen ('-') instead. This may lead to multiple classes the process could be assigned to, depending on the values of the unspecified attributes. If all the attributes are left unspecified, an error is returned.

The *attributes* string is provided in a format identical to the format of the attributes in the rules file: a list of attribute values separated by spaces. The order of the attributes in the assignment rules is:

1. reserved: must be a hyphen ('-')
2. user name
3. group name
4. application pathname
5. type of application
6. tag

Each field can have at most one value. Exclusion (!), attribute value groupings (\$), comma separated lists and wild cards are not allowed. For the type field, the AND operator "+" is allowed, since a process can have several of the possible values for the type attribute at the same time. For instance a process can be a 32 bit process and call plock, or be a 64 bit fixed priority process.

Here are examples of valid *attributes* strings:

```
"- bob staff /usr/bin/emacs - -"
"- - - /usr/sbin/dbserve - _DB1"
"- - devlt - 32bit+fixed"
"- sally"
```

The class name(s) returned by the function in the *class* buffer is fully-qualified, null-terminated class names of the form **supername.subname**.

This function does not require any special privileges and can be called by all users.

## Parameters

<i>config</i>	Specifies a pointer to a string containing the name of a valid Workload Manager (WLM) configuration (the name of a subdirectory of <i>/etc/wlm</i> ). If a null string ( <code>'\0'</code> ) is given, the <b>wlm_classify</b> subroutine uses <i>current</i> as the default configuration.  If the configuration is a set of time-based configurations, either because <i>config</i> or <i>current</i> is a configuration set, the subroutine will apply to the currently applicable configurations of the set.
<i>attributes</i>	Specifies the address of a string, with the format described above, containing a list of values for the process attributes used for automatic classification of processes.
<i>class</i>	Specifies a pointer to a buffer where the name of the class the process could be assigned to is returned as consecutive null-terminated character strings.
<i>len</i>	Specifies a pointer to an integer containing the length in bytes of the buffer pointed to by <i>class</i> when calling <b>wlm_classify</b> and the actual number of class names copied into the <i>class</i> buffer upon successful return.

## Return Values

Upon successful completion, the **wlm\_classify** subroutine returns a value of 0. In case of error, a non-0 value is returned.

When a non-0 value is returned, the content of the **class** buffer and the value of the integer pointed to by **len** are unspecified.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the **wlm.h** header file.

## Related Information

The **wlmcheck** command.

The **wlm.h** header file.

Workload Manager rules File in *AIX 5L Version 5.2 Files Reference*.

Automatic assignment (“wlm\_classify Subroutine” on page 477) in *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

---

## wlm\_class2key Subroutine

### Purpose

Class name to key translation.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

```
int wlm_class2key ( struct wlm_args *args, wlm_key_t *key)
```

## Description

The **wlm\_class2key** subroutine generates a 64-bit numeric key from a WLM class name. The **wlm\_class2key** subroutine is provided for applications gathering high volumes of per-class usage statistics or accounting data and allows those applications to save storage space by compressing the class name (up to 34 characters long) into a 64-bit integer. The **wlm\_key2class** subroutine can then get the key-to-class name conversion for data reporting purposes

## Parameters

*wlm\_args* Only 2 fields need to be initialized in the **wlm\_args** structure pointed to by **args**:

- *cl\_def.data.descr.name* specifies the null terminated full name of the class (<super\_name>.<subname> for a subclass).
- *versflags* initialized with WLM\_VERSION and optionally WLM\_MUTE.

## Return Values

If the **wlm\_class2key** subroutine is successful, a value of 0 is returned. If the **wlm\_class2key** subroutine is unsuccessful, an error code is returned.

## Error Codes

If the **wlm\_class2key** subroutine is unsuccessful, one of the following error codes is returned:

<i>WLM_NOT_INITED</i>	Missing call to <b>wlm_init</b> .
<i>WLM_EFAULT</i>	Invalid key or args pointer.
<i>WLM_BADCNAME</i>	The class name contains invalid characters.

## Related Information

The **wlm\_endkey** subroutine.

The **wlm\_initkey** subroutine.

The **wlm\_key2class** subroutine.

---

## wlm\_create\_class Subroutine

### Purpose

Creates a new Workload Manager (WLM) class.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

```
int wlm_create_class ( wlmargs )
```

```
struct wlm_args *wlmargs;
```

### Description

The **wlm\_create\_class** subroutine creates a new class for a given WLM configuration using the values passed in the data structure of type **struct wlm\_args** pointed to by *wlmargs*.

- If the name of a configuration is passed in the **confdir** field, the subroutine updates the WLM properties files for the target configuration. When creating the first subclass of a superclass, the subroutine creates a subdirectory of **/etc/wlm/<confdir>** with the name of the superclass and create the WLM properties files in this new directory. The newly created properties files have entries for the Default and Shared subclass automatically created in addition to entries for the new subclass.
- If a null string ('') is passed in the **confdir** field, the new superclass or subclass is created only in the in-core WLM data. No WLM properties file are updated. In that case, the new class definition is lost if WLM is stopped and restarted, or if the system reboots.

The structure of type **struct class\_definition**, which is part of **struct wlm\_args**, has normally been initialized with a call to **wlm\_init\_class\_definition**. Once this has been done, initialize the fields of this structure which have no default value (such as the name of the new class) or for which the desired value is different from the default value. For a description of the possible values for all the class attributes and their default values, refer to the description of **wlm.h** in the *AIX 5L Version 5.2 Files Reference*.

The caller must have root authority to create a superclass and must have administrator authority on a superclass to create a subclass of the superclass.

**Note:** Do not specify a set in the *confdir* field of the **wlm\_args** structure. The **wlm\_create\_class** subroutine cannot apply to a set of time-based configurations.

## Parameter

**wlmargs** Specifies the address of the **struct wlm\_args** data structure containing the **class\_definition** structure for the new class to be created.

The following fields of the **wlm\_args** structure and the embedded substructures need to be provided:

<b>versflags</b>	Needs to be initialized with <b>WLM_VERSION</b> .
<b>confdir</b>	Specifies the name of the WLM configuration the new class is to be added to. It must be either the name of a valid subdirectory of <b>/etc/wlm</b> or an empty string (starting with '').  If the name is a valid subdirectory, the new class data is added to the given WLM configuration's class description files.  If the name is a null string, no description files are updated. The new class is created and the data is passed to the kernel immediately.
<b>name</b>	Specifies the name of the superclass or of the subclass to be created. If this is a subclass name, it must be of the form <b>super_name.sub_name</b> . There is no default for this field.

All the other fields can be left at their default value if the user does not wish to use specific values.

## Return Values

Upon successful completion, the **wlm\_create\_class** subroutine returns a value of 0. If the **wlm\_create\_class** subroutine is unsuccessful, a nonzero value is returned.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the **wlm.h** header file.

## Related Information

The **mkclass** command, **chclass** command, **rmclass** command.

The **wlm.h** header file.

The **wlm\_change\_class** (“wlm\_change\_class Subroutine” on page 475) subroutine, **wlm\_delete\_class** (“wlm\_delete\_class Subroutine”) subroutine.

Workload Management in *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

---

## wlm\_delete\_class Subroutine

### Purpose

Deletes a class.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

```
int wlm_delete_class ( wlmargs)
```

```
struct wlm_args *wlmargs;
```

### Description

The **wlm\_delete\_class** subroutine deletes an existing superclass or subclass. A superclass cannot be deleted if it still has subclasses other than Default and Shared defined.

- If the name of a valid configuration is passed in the **confdir** field, the subroutine updates the Workload Manager (WLM) properties files for the target configuration, removing all references to the class to be deleted.
- If a null string ('\0') is passed in the **confdir** field, the class is deleted only from the in-core WLM data structures. No WLM properties file is updated. This is normally used to delete a class which was also only created in the in-core WLM data structures. Otherwise, the class deletion is temporary and the class will be created again when WLM is updated or restarted with a configuration where the class exists in the classes file.

The caller must have root authority to delete a superclass and must have administrator authority on a superclass to delete a subclass of the superclass.

**Note:** Do not specify a set in the *confdir* field of the **wlm\_args** structure. The **wlm\_delete\_class** subroutine cannot apply to a set of time-based configurations.

### Parameter

*wlmargs*

Specifies the address of the **struct wlm\_args** data structure containing the information about the class to be deleted.

The following fields of the **wlm\_args** structure and the embedded substructures need to be provided:

**versflags**

Needs to be initialized with **WLM\_VERSION**.

**confdir**

Specifies the name of the WLM configuration the target class belongs to. It must be either the name of a valid subdirectory of **/etc/wlm** or an empty string (starting with `'\0'`).

If the name is a valid subdirectory, the relevant class description files in the specified configuration are modified.

If the name is a null string, no description files are updated. The class is removed from the kernel WLM data structures.

**name**

Specifies the name of the superclass or of the subclass to be deleted. If this is a subclass name, it must be of the form **super\_name.sub\_name**. There is no default for this field.

All the other fields can be left uninitialized for this call.

## Return Values

Upon successful completion, the **wlm\_delete\_class** subroutine returns a value of 0. If the **wlm\_delete\_class** subroutine is unsuccessful, a non-0 value is returned.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the **wlm.h** header file.

## Related Information

The **mkclass** command, **chclass** command, **rmclass** command.

The **wlm.h** header file.

The **wlm\_change\_class** (“wlm\_change\_class Subroutine” on page 475) subroutine, **wlm\_create\_class** (“wlm\_create\_class Subroutine” on page 480) subroutine.

Workload Management in *AIX 5L Version 5.2 System Management Concepts: Operating System and Devices*.

---

## wlm\_endkey Subroutine

### Purpose

Frees the classes to keys translation table.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

```
int wlm_endkey(struct wlm_args *args, void *ctx)
```

## Description

The **wlm\_endkey** subroutine frees the classes to the keys translation table. The memory area pointed to by *ctx* is freed.

## Parameters

- *ctx* Points to the memory area to be freed.  
*wlm\_args* A pointer to a **wlm\_args** structure:

**versflag** field is the only field in the structure that needs to be initialized with WLM\_VERSION and optionally WLM\_MUTE.

## Return Values

When the **wlm\_endkey** operation is successful, it returns a value of 0, and if it is unsuccessful, it returns an error code.

## Error Codes

If the **wlm\_endkey** subroutine is unsuccessful, one of the following error codes is returned:

<i>WLM_BADVERS</i>	Bad version number.
<i>WLM_NOT_INITED</i>	Missing call to <b>wlm_init</b> .
<i>WLM_EFAULT</i>	Invalid <i>ctx</i> or <i>args</i> argument.

## Related Information

The **wlm\_class2key** subroutine.

The **wlm\_initkey** subroutine.

The **wlm\_key2class** subroutine.

---

## wlm\_get\_bio\_stats subroutine

### Purpose

Read the WLM disk I/O statistics per class or per device.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

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

```
int wlm_get_bio_stats ( dev, array, count, class, flags)
```

```
dev_t dev;
```

```
void *array;
```

```
int *count;
```

```
char *class;
```

```
int flags;
```

## Description

The `wlm_get_bio_stats` subroutine is used to get the WLM disk IO statistics. There are two types of statistics available:

- The statistics about disk IO utilization per class and per devices, returned by `wlm_get_bio_stats` in `wlm_bio_class_info_t` structures,
- The statistics about the disk IO utilization per device, all classes combined, returned by `wlm_get_bio_stats` in `wlm_bio_dev_info_t` structures.

The type of statistics returned by the function is predicated on the value of the `flags` argument. The `flags` argument, together with the `dev` and `class` arguments, are used to restrict the scope of the function to a class or a set of classes and/or a device or a set of devices. If the value passed to the routine in the `count` argument is equal to zero (0), `wlm_get_bio_stats` does not copy any device statistics (and, in this case, the `array` argument can be a NULL pointer but sets this count to the number of elements in scope for the specific set of parameters. This is a way of finding out how big an array is needed to get all the information for a given set of classes and devices.

`wlm_get_bio_stats` does not require any special privileges and is accessible to all users.

`wlm_get_bio_stats` fails if WLM is off.

## Parameters

`flags`

Need to be initialized with **WLM\_VERSION**. Optionally, the following flag values can be or'ed to **WLM\_VERSION**:

**WLM\_SUPER\_ONLY**

Limits the scope to superclasses only

**WLM\_SUB\_ONLY**

Limits the scope to subclasses only

**WLM\_BIO\_CLASS\_INFO**

Per class statistics requested

**WLM\_BIO\_DEV\_INFO**

Per device statistics requested

**WLM\_BIO\_ALL\_DEV**

Requests statistics for all devices. When this flag is set, the value passed in the `dev` argument is ignored.

**WLM\_BIO\_ALL\_MINOR**

Requests statistics for all devices associated with a given major number. When this flag is set, only the major number part of the value passed in the `dev` argument is used.

**WLM\_VERBOSE\_MODE**

Shows the system defined subclasses (*Default* and *Shared*) even if they have not been modified by a WLM administrator.

One of the flags **WLM\_BIO\_CLASS\_INFO** or **WLM\_BIO\_DEV\_INFO** (and only one) must be specified. **WLM\_SUPER\_ONLY** and **WLM\_SUB\_ONLY** are mutually exclusive.

<i>dev</i>	Device identification (major, minor) of a disk device. <ul style="list-style-type: none"> <li>• If <i>dev</i> is equal to 0, the statistics for all devices are returned (even if <b>WLM_BIO_ALL_DEV</b> is not specified in the <i>flags</i> argument).</li> <li>• If <i>dev</i> is not equal to 0 and <b>WLM_BIO_ALL_MINOR</b> is specified in the <i>flags</i> argument, the statistics for all disk devices with the same major number specified in <i>dev</i> are returned.</li> <li>• If <i>dev</i> is not equal to 0 and <b>WLM_BIO_ALL_MINOR</b> is not specified in the <i>flags</i> argument, only the statistics for the disk device with the major and minor numbers specified in <i>dev</i> are returned.</li> </ul>
<i>array</i>	Pointer to an array of <b>wlm_bio_class_info_t</b> structures (when <b>WLM_BIO_CLASS_INFO</b> is specified in the <i>flags</i> argument) or an array of <b>wlm_bio_dev_info_t</b> structures (when <b>WLM_BIO_DEV_INFO</b> is specified in the <i>flags</i> argument). A NULL pointer can be passed together with a <i>count</i> of 0 to determine how many elements are in scope for the set of arguments passed.
<i>count</i>	The address of an integer containing the maximum number of elements to be copied into the array above. If the call to <b>wlm_get_bio_stats</b> is successful, this integer will contain the number of elements actually copied. If the initial value is equal to zero (0), <b>wlm_get_bio_stats</b> sets this value to the number elements selected by the specified combination of flags and class.
<i>class</i>	A pointer to a character string containing the name of a superclass or subclass. If <i>class</i> is a pointer to an empty string (""), the information for all classes are returned. The <i>class</i> parameter is taken into account only when the flag <b>WLM_BIO_CLASS_INFO</b> is set.

## Return Values

Upon successful completion, a value of 0 is returned and the value pointed to by *count* is set to the number of elements copied into the array of structures pointed to by *array*. If the **wlm\_get\_bio\_stats** subroutine is unsuccessful a non 0 value is returned.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the header file **sys/wlm.h**.

## Related Information

The **wlm.h** header file.

## wlm\_get\_info Subroutine

### Purpose

Read the characteristics of superclasses or subclasses.

### Library

Workload Manager Library (**libwlm.a**)

## Syntax

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

int wlm_get_info ( wlmargs, info, count)

struct wlm_args *wlmargs;

struct wlm_info *info

int *count
```

## Description

The **wlm\_get\_info** subroutine is used to get the characteristics of the classes defined in the active Workload Manager (WLM) configuration, together with their current resource usage statistics. For a detailed description of the fields of the structure **wlm\_info**, refer to the description of the **wlm.h** header file in the *AIX 5L Version 5.2 Files Reference* documentation.

By default, the scope of the **wlm\_get\_info** subroutine is all the superclasses and all the subclasses. This scope can be limited to a subset of the classes using flags in the **versflags** field of **wlm\_args** or a superclass or subclass name in the **name** field of the substructure **class\_definition** of **wlm\_args**.

The information related to the superclasses and subclasses within the scope of **wlm\_get\_info** are copied to the array of **wlm\_info** structures pointed to by **info**. The total number of classes for which information is copied to the array at **info** is limited to the value of the integer pointed to by **count**. If the routine is successful, the value of the integer pointed to by **count** is set to the actual number of classes copied. If the value passed to the routine for the count is equal to zero (0), **wlm\_get\_info** does not copy any class statistics but sets this count to the number of classes in scope for the specific set of parameters. This is a way of finding out how big an array is needed to get all the information for a given set of classes (superclasses or subclasses).

This is a way of finding out how big an array is needed to get all the information for a given set of classes (superclasses or subclasses).

The **wlm\_get\_info** subroutine does not require any special privileges and is accessible to all users. **wlm\_get\_info** fails if WLM is off.

## Parameters

### *wlmargs*

The address of a **struct wlm\_args** data structure.

The following fields of the **wlm\_args** structure and the embedded substructures need to be provided:

### **versflags**

Needs to be initialized with **WLM\_VERSION**. Optionally, the following flag value can be or'ed to **WLM\_VERSION**:

#### **WLM\_SUPER\_ONLY**

Limits the scope to superclasses only

#### **WLM\_SUB\_ONLY**

Limits the scope to subclasses only

#### **WLM\_VERBOSE\_MODE**

Shows the system-defined subclasses (Default and Shared) even if they have not been modified by a WLM administrator.

**WLM\_SUPER\_ONLY** and **WLM\_SUB\_ONLY** are mutually exclusive.

- name** Contains either a null string or the name of a valid superclass or subclass (in the form **Super.Sub**). This field can be used in conjunction with the flags to further narrow the scope of **wlm\_get\_info**:
- If the name of a subclass is provided, **wlm\_get\_info** returns the statistics only for the specified subclass.
  - If the name of a superclass is provided or if none of the **WLM\_SUPER\_ONLY** and **WLM\_SUB\_ONLY** flag is provided, **wlm\_get\_info** returns the statistics for the specified superclass and all its subclasses.
  - If the name of a superclass is provided together with **WLM\_SUPER\_ONLY**, **wlm\_get\_info** returns only the statistics for the specified superclass.
  - If the name of a superclass is provided together with **WLM\_SUB\_ONLY**, **wlm\_get\_info** returns the statistics for all the subclasses of the specified superclass.

All the other fields of the **wlm\_args** structure can be left uninitialized.

*info* The address of an array of structures of type **struct wlm\_info**. Upon successful return from **wlm\_get\_info**, this array contains the WLM statistics for the classes selected.

*count* The address of an integer containing the maximum number of element (of type **wlm\_info**) for **wlm\_get\_info** to copy into the array above. If the call to **wlm\_get\_info** is successful, this integer contains the number of elements actually copied. If the initial value is equal to zero (0), **wlm\_get\_info** sets this value to the number of classes selected by the specified combination of **versflags** and **name** above.

## Return Values

Upon successful completion, the **wlm\_get\_info** subroutine returns a value of 0. If the **wlm\_get\_info** subroutine is unsuccessful a non-0 value is returned.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the **wlm.h** header file.

## Related Information

The **wlmstat** command.

The **wlm.h** header file.

---

## wlm\_get\_procinfo Subroutine

### Purpose

Retrieves per-process Workload Manager information.

### Library

Workload Manager Library (**libwlm.a**)

## Syntax

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

int wlm_get_procinfo (pid, wlm_pinfo)
pid_t pid;
struct wlm_procinfo *wlm_pinfo;
```

## Description

The **wlm\_get\_procinfo** subroutine returns Workload Manager information for the process associated with the *pid* parameter, into the buffer pointed to by the *wlm\_pinfo* parameter. If process total accounting is disabled, the related fields (*totalconnecttime*, *termtime*, *totalcputime*, and *totaldiskio*) are set to -1. When WLM is on, the class name of the process is set in the *classname* field of the **wlm\_procinfo** structure. When WLM is off, this field is set to *Unclassified*.

## Parameters

<i>pid</i>	Indicates from which process to retrieve the Workload Manager information.
<i>wlm_pinfo</i>	Points to the buffer where the Workload Manager information is stored.

## Return Values

Upon successful completion, the **wlm\_get\_procinfo** subroutine returns a zero. If the **wlm\_get\_procinfo** subroutine is unsuccessful, a nonzero value is returned.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the **wlm.h** header file.

## Related Information

The **wlm.h** header file.

---

## wlm\_init\_class\_definition Subroutine

### Purpose

Initializes a variable of type **struct class\_definition**, defined in **<sys/wlm.h>** for use as an argument to Workload Manager (WLM) API function calls.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

int wlm_init_class_definition ( wlmargs)

struct wlm_args *wlmargs;
```

### Description

The **wlm\_init\_class\_definition** subroutine initializes or reinitializes the data structure of type **struct class\_definition**, which is part of the argument of type **struct wlm\_args** pointed to by *wlmargs* (field **class**), so that this data structure can be used as an argument for the class management subroutines of the WLM API library. The purpose of this call is to allow applications to initialize only the fields that are

relevant for the operation they execute. For example, to change a CPU limit or share for an existing class after a call to **wlm\_init\_class\_definition**, the application has to initialize the fields corresponding to the values it wishes to modify.

This routine initializes all values to specific invalid values so that the WLM library routines can find out which fields have been explicitly initialized by the user. This way, they can set or modify only the corresponding attributes. When creating a class, for instance, it is different to leave a **class** attribute at its invalid value set by **wlm\_initialize** than setting its value to the current default value for the attribute. In the former case, the attribute will not appear in the property file. In the latter, it will appear and will be set with the value passed.

This makes a difference if a WLM administrator decides to change the default value for an attribute using the special stanza default in a property file. For instance, the system default for the **inheritance** attribute is no. If a WLM administrator wants the inheritance to be yes by default, using this special stanza, all the classes in the classes property file, for which the **inheritance** attribute has not been specified, will now use the default of yes. Those for which the **inheritance** attribute has been specified with its old default of no will not have inheritance.

## Parameter

*wlmargs*

Specifies the address of the **struct wlm\_args** data structure containing the **class\_definition** structure to be initialized.

Only the **versflags** field of the **wlm\_args** structure passed need to be initialized with **WLM\_VERSION**.

## Return Values

Upon successful completion, the **wlm\_init\_class\_definition** subroutine returns a value of 0. If the **wlm\_init\_class\_definition** subroutine is unsuccessful a non-0 value is returned.

## Error Codes

There are two possible error code returned by **wlm\_init\_class\_definition**:

**BADVERSION**

Specifies the value of the flags parameter is not a supported version number.

**NOTINITED**

Specifies the WLM API has not been initialized by a prior call to **wlm\_init**.

## Related Information

The **wlm.h** header file.

The **wlm\_change\_class** (“wlm\_change\_class Subroutine” on page 475) subroutine, **wlm\_create\_class** (“wlm\_create\_class Subroutine” on page 480) subroutine, **wlm\_delete\_class** (“wlm\_delete\_class Subroutine” on page 482) subroutine.

---

## wlm\_initialize Subroutine

### Purpose

Prepares Workload Manager (WLM) for use by an application.

### Library

Workload Manager Library (**libwlm.a**)

## Syntax

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

```
int wlm_initialize ( flags)
```

```
int flags;
```

## Description

The **wlm\_initialize** subroutine initializes the WLM API for use with an application program. It is mandatory to call **wlm\_initialize** prior to using the WLM API. Otherwise, all other WLM API function calls return an error.

## Parameter

*flags*

Specifies that the format is the same as the **versflag** field of the **wlm\_args** structure. The value for the argument must have the version number in the upper 4 bits (**WLM\_VERSION**) possibly or'ed with a flag in the lower 28 bits.

## Return Values

Upon successful completion, the **wlm\_initialize** subroutine returns a value of 0. If the **wlm\_initialize** subroutine is unsuccessful a non-0 value is returned.

## Error Codes

There are two possible error codes returned by **wlm\_initialize**:

**BADVERSION**

The value of the *flags* parameter is not a supported version number.

**WLMINITED**

There has already been a previous call to **wlm\_initialize**.

## Related Information

The **wlm.h** header file.

---

## wlm\_initkey Subroutine

### Purpose

Allocates and initializes the classes to keys translation table.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

```
int wlm_initkey ( struct wlm_args *args, void **ctx)
```

### Description

The **wlm\_initkey** subroutine allocates a block of memory, builds the keys <==> class names translation table and returns its address into the **ctx** argument.

## Parameters

*args*

Only 2 fields need to be initialized in the **wlm\_args** structure pointed to by **args**:

- *confdir* specifies the null-terminated name of the WLM configuration to be searched (the name can be "current" to specify the current configuration). If the configuration name passed is an empty string (starts with '\0'), then all the configurations in **/etc/wlm** are searched.
- *versflags* initialized with WLM\_VERSION and optionally WLM\_MUTE.

## Return Values

If the **wlm\_initkey** subroutine is successful, a value of 0 is returned. If the **wlm\_initkey** subroutine is unsuccessful, an error code is returned.

## Error Codes

If the **wlm\_initkey** subroutine is unsuccessful, one of the following error codes is returned:

<i>WLM_BADVERS</i>	Bad version number.
<i>WLM_NOT_INITED</i>	Missing call to <b>wlm_init</b> .
<i>WLM_NOMEM</i>	Not enough memory.
<i>WLM_NOCLASS</i>	Specified configuration does not exist.
<i>WLM_EFAULT</i>	Invalid <b>ctx</b> or <b>args</b> argument.

## Related Information

The **wlm\_endkey** subroutine.

The **wlm\_class2key** subroutine.

The **wlm\_key2class** subroutine.

---

## wlm\_key2class Subroutine

### Purpose

Retrieves a class name from a key.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

```
int wlm_key2class ( struct wlm_args *args, wlm_key_t key, void *ctx)
```

### Description

The **wlm\_key2class** subroutine retrieves a class name from a 64-bit key calculated using the **wlm\_class2key** subroutine. The key-to-class translation is made by going through the WLM configuration files for the configuration named in the **wlm\_args** structure pointed to by **args** (or all the WLM configuration files, if no configuration name is given), and translating all the class names to a 64-bit key until the matching key is found.

This process is time consuming and WLM offers the subroutines **wlm\_initkey** and **wlm\_endkey** for applications needing to translate several 64-bit keys back to class names. These subroutines can be used in conjunction with the **wlm\_key2class** subroutine to speed up searches.

The **wlm\_initkey** subroutine allocates a block of memory, calculates the keys corresponding to the class names in the configuration(s) in scope, stores the names with the corresponding keys in the memory buffer, and returns its address. This address is passed to the **wlm\_key2class** subroutine using the **ctx** argument, so that **wlm\_key2class** only needs to search through the memory buffer.

After all keys have been translated into class names, the application must call **wlm\_endkey** to free the memory buffer. Alternatively, for an application translating only one key, it is possible to call **wlm\_key2class** directly using a null pointer in the **ctx** argument. This causes the **wlm\_key2class** subroutine to internally call **wlm\_initkey** and **wlm\_endkey**.

The method of retrieving class names through the WLM configuration files implies that if a class has been deleted between the time the class name was converted into a key and the call to the **wlm\_key2class** subroutine, the name corresponding to the key will not be found and the **wlm\_key2class** subroutine returns an error.

## Parameters

- *args*                      A pointer to a **wlm\_args** structure:
  - **confdir** field needs to be initialized as described in **wlm\_initkey** if **wlm\_initkey** has not been previously invoked (**ctx** == NULL). Otherwise, the **confdir** field is ignored.
  - **versflags** field needs to be initialized with WLM\_VERSION and optionally WLM\_MUTE.
- *ctx*                        The context handler returned by **wlm\_initkey**, or a NULL pointer otherwise. .
- *key*                        The search key.

## Return Values

When the **wlm\_key2class** operation is successful, the first class name matching the value of the key is returned in the name sub-field of the **wlm\_args** structure pointed to by **args**.

## Error Codes

If the **wlm\_key2class** subroutine is unsuccessful, one of the following error codes is returned:

<i>WLM_BADVERS</i>	Bad version number.
<i>WLM_NOT_INITED</i>	Missing call to <b>wlm_init</b> .
<i>WLM_NOMEM</i>	Not enough memory.
<i>WLM_NOCLASS</i>	No class matching the key was found.
<i>WLM_EFAULT</i>	Invalid <i>ctx</i> or <i>args</i> argument.

## Related Information

The **wlm\_class2key** subroutine.

The **wlm\_endkey** subroutine.

The **wlm\_initkey** subroutine.

---

## wlm\_load Subroutine

### Purpose

Loads a Workload Manager (WLM) configuration into the kernel.

## Library

Workload Manager Library (**libwlm.a**)

## Syntax

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

```
int wlm_load ( wlmargs )
```

```
struct wlm_args *wlmargs;
```

## Description

The **wlm\_load** subroutine loads into the kernel the property files for the WLM configuration passed in the *confdir* field of the **wlmargs** structure. The *confdir* field may also refer to a set of time-based configurations, in which case the appropriate configuration of the set will be loaded and the WLM daemon will later switch to the other configurations of the set on a time basis.

If the WLM is running and *confdir* is not current, this leads to switch to the specified configuration (or configuration set).

If the WLM is running and *confdir* is current, **wlm\_load** will refresh the current WLM configuration into the kernel. If a superclass name is given in the *name* field of the *class\_definition* substructure, only the subclasses of the given superclass are refreshed. In this context:

- The **wlm\_load** subroutine is accessible to root users and to users with administration privileges on the subclasses of the superclass. In all other cases, the **wlm\_load** subroutine is only accessible to root users.
- The **wlm\_load** subroutine cannot be used to change the mode of operation of WLM (for example, to switch between active and passive modes).
- If *current* is a configuration set, *confdir* must be given in the form *current/config* where *config* is the regular configuration of the set the superclass belongs to. If *config* is the active configuration of the set, the changes will take effect immediately, otherwise they will take effect the next time *config* is made active.

If the caller of **wlm\_load** has root privileges and does not specify a superclass, the flags passed in *versflags* can be used to start WLM in active or passive mode, switch between active and passive modes, or enable/disable the rset bindings or the process or class total limits. The **wlm\_load** subroutine cannot be used to stop WLM. Use the **wlm\_set** subroutine instead.

## Parameter

*wlmargs* Specifies the address of the **struct wlm\_args** data structure containing information about the configuration (or configuration set or superclass) to be loaded and the mode of operation of WLM.

The following fields of the **wlm\_args** structure and the embedded substructures can be provided:

**versflags** Needs to be initialized with WLM\_VERSION. May be ORed with WLM\_MUTE for **wlm\_load** to be silent.

If no change must be done to the mode of operation of WLM, it must be ORed with WLM\_TEST\_ON (mandatory if superclass is specified).

Otherwise, one of the mutually exclusive flags (WLM\_ACTIVE, WLM\_CPUONLY, or WLM\_PASSIVE) must be given. One or more of the WLM\_BIND\_RSETS, WLM\_PROCTOTAL, or WLM\_CLASSTOTAL flags can be given optionally.

<b>confdir</b>	Specifies the name of the WLM configuration to be loaded into the kernel. It must be either the name of a valid configuration or configuration set in the <b>/etc/wlm</b> subdirectory, the <i>current</i> string to refer to the active configuration, or, if superclass is specified and current is a configuration set, it must indicate which configuration of current set the superclass belongs to in the form: <i>current/config</i> (this is different from specifying <i>config</i> only, which is considered a configuration switch request).
<b>name</b>	Specifies the name of a <b>superclass</b> . This is used to refresh only the subclasses of a given superclass.

## Return Values

Upon successful completion, the **wlm\_load** subroutine returns a value of 0. If the **wlm\_load** subroutine is unsuccessful, a nonzero value is returned.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the **wlm.h** header file.

## Related Information

The **wlmcntrl** command.

The **wlm\_set** (“wlm\_set Subroutine” on page 497) subroutine.

The **wlm.h** header file.

## wlm\_read\_classes Subroutine

### Purpose

Reads the characteristics of superclasses or subclasses.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

int wlm_read_classes (wlmargs, class_tbl, nclass)
struct wlm_args *wlmargs;
struct class_definition *class_tbl;
int *nclass;
```

### Description

The **wlm\_read\_classes** subroutine is used to get the characteristics of the superclasses or the subclasses of a given subclass of a Workload Manager (WLM) configuration.

- If the name of a configuration is passed in the **confdir** field, the **wlm\_read\_classes** subroutine reads the property files of the classes of the specified configuration. If **confdir** is set to a null string (‘\0’), **wlm\_read\_classes** reads the classes’ characteristics from the in-core WLM data structures when WLM is on (and returns an error when WLM is off).

**Note:** These values may be different from the values in the property files of the configuration pointed to by **/etc/wlm/current**. For instance when a WLM administrator has modified the property files for

the configuration pointed to by **/etc/wlm/current** but has not refreshed WLM yet. Another example is if applications dynamically created or modified classes through the API without saving the changes in the *current* configuration property files.

If your application specifically needs to access the properties of the classes as described in the **/etc/wlm/current** configuration, you must specify *current* as the configuration name in **confdir**.

If the name of a set of time-based configurations is passed in the *confdir* field, the **wlm\_read\_classes** subroutine reads the classes of the currently applicable configuration of the set.

- If the name of a valid superclass of the given configuration is passed in the **name** field of the **class\_descr** substructure of *wlmargs*, **wlm\_read\_classes** reads the property files for the subclasses of this superclass. If a null string ('') is passed in the **name** field, **wlm\_read\_classes** reads the property files for the superclasses of the WLM configuration described above.
- When **wlm\_read\_classes** is successful, the characteristics of the superclasses or subclasses are copied into the array of **class\_definition** structures pointed to by *class\_tbl*. The integer value pointed to by *nclass* indicates the maximum number of class definitions to be copied. Upon successful return from the function, this value reflects the actual number of classes read.

If the number of elements copied by **wlm\_read\_classes** is strictly smaller than the number of elements passed as an argument, all the classes have been read. If it is equal, it may mean that some classes were not copied into the **class\_tbl** array because its size is too small.

The maximum number of classes read by **wlm\_read\_classes** is 67 (64 user-defined superclasses plus System, Shared and Default) when reading superclasses and 63 (61 user-defined subclasses plus Shared and Default) when reading subclasses characteristics.

- Upon successful return from **wlm\_read\_classes**, the substructure **class** of type **struct class\_definition** of the structure pointed to by *wlmargs* contains the default values of various class attributes for the returned set of classes.

This operation does not require any special privileges and is accessible to all users.

## Parameter

<i>wlmargs</i>	<p>Specifies the address of a <b>struct wlm_args</b> data structure.</p> <p>The following fields of the <b>wlm_args</b> structure and the embedded substructures need to be provided:</p> <p><b>versflags</b> Needs to be initialized with <b>WLM_VERSION</b>.</p> <p><b>confdir</b> Specifies the name of a WLM configuration. It must be either the name of a valid subdirectory of <b>/etc/wlm</b> or a null string (starting with '').</p> <p><b>name</b> Specifies the name of a superclass existing in the specified configuration or a null string.</p> <p>All the other fields can be left uninitialized.</p>
<i>class_tbl</i>	<p>Specifies the address of an array of structures of type <b>struct class_definition</b>. Upon successful return from <b>wlm_read_classes</b>, this array contains the characteristics of the classes read.</p>
<i>nclass</i>	<p>Specifies the address of an integer containing the maximum number of element (class definitions) for <b>wlm_read_classes</b> to copy into the array above. If the call to <b>wlm_read_classes</b> is successful, this integer contains the number of elements actually copied.</p>

## Return Values

Upon successful completion, the `wlm_read_classes` subroutine returns a value of 0. If the `wlm_read_classes` subroutine is unsuccessful, a nonzero value is returned.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the `wlm.h` header file.

## Related Information

The `Isclass` command.

The `wlm.h` header file.

---

## wlm\_set Subroutine

### Purpose

Sets or queries the Workload Manager (WLM) state.

### Library

Workload Manager Library (`libwlm.a`)

### Syntax

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

```
int wlm_set ( flags )  
int *flags;
```

### Description

The `wlm_set` subroutine is used to set, change, or query the mode of operations of WLM. The state of WLM can be:

<b>OFF</b>	Does not classify processes, monitor or regulate resource utilization.
<b>ON</b> in <b>passive</b> mode	Classifies the processes and monitors their resource usage but does no regulation.
<b>ON</b> in <b>active</b> mode	Specifies the normal operating mode where WLM classifies processes, monitors and regulates the resource usage.

## Parameters

- flags* Specifies the address of an integer interpreted in a manner similar to the **versflags** field of the **wlmargs** structure passed to the other API routines. The integer pointed to by *flags* should be initialized with **WLM\_VERSION**. In addition, one or more of the following values can be or'ed to **WLM\_VERSION**:
- WLM\_TEST\_ON**  
Queries the state of WLM without altering it.
  - WLM\_OFF**  
Turns WLM off.
  - WLM\_ACTIVE**  
Turns WLM on in **active** mode or transitions from any mode to **active** mode.
  - WLM\_CPU\_ONLY**  
Turns WLM on in **active** mode for CPU resource only, or transitions from any mode to this mode. This is the same as **WLM\_ACTIVE**, but only CPU resources are regulated. Other resources (memory, disk IO, and total limits when enabled) are still accounted.
  - WLM\_PASSIVE**  
Turns WLM on in **passive** mode or transitions from any mode to **passive** mode.
  - WLM\_BIND\_RSETS**  
Requests that WLM takes the resource set bindings into account.
  - WLM\_PROCTOTAL**  
Enables process total limits on resource usage.
  - WLM\_CLASSTOTAL**  
Enables class total limits on resource usage.

Some combinations of the flags above are not legal:

- **WLM\_OFF**, **WLM\_ACTIVE**, **WLM\_CPU\_ONLY**, and **WLM\_PASSIVE** are mutually exclusive.
- **WLM\_BIND\_RSETS**, **WLM\_PROCTOTAL**, and **WLM\_CLASSTOTAL**, are ineffective when used together with **WLM\_OFF**.
- Only **WLM\_TEST\_ON** is allowed to non-root users.
- If **WLM\_TEST\_ON** is specified, the other flags are ineffective and should not be specified.

## Return Values

Upon successful completion, the **wlm\_set** subroutine returns a value of 0, and the current state of WLM is returned in the *flags* parameter. The return value is **WLM\_OFF**, **WLM\_ACTIVE**, **WLM\_CPU\_ONLY**, or **WLM\_PASSIVE**. When WLM is on in either mode, the **WLM\_BIND\_RSETS**, **WLM\_PROCTOTAL**, and **WLM\_CLASSTOTAL**, flags are added when appropriate.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the **wlm.h** header file.

## Related Information

The **wlmcntrl** command.

The **wlm.h** header file.

The **wlm\_load** (“wlm\_load Subroutine” on page 493) subroutine.

---

## wlm\_set\_tag Subroutine

### Purpose

Sets the current process's tag and related flags.

### Library

Workload Manager Library (**libwlm.a**)

### Syntax

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

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

```
int wlm_set_tag ( tag, flags)
```

```
char *tag;
```

```
int *flags;
```

### Description

The **tag** attribute is a new attribute of a process that can be set using the Workload Manager (WLM) **wlm\_set\_tag** subroutine. This tag is a character string with a maximum length of **WLM\_TAG\_LENGTH** (not including the null terminator). Process tags can be displayed using the **ps** command.

The **tag** attribute is also one of the **process** attributes used in the assignment rules to automatically assign a process to a given class. The syntax of the assignment rules precludes the use of special characters in the application tag string. Thus, application tags should be comprised only of upper and lower case letters, numbers and underscores ('\_').

The main use of the **tag** attribute is to allow WLM administrators to discriminate between several instances of the same application, which typically have the same user and group ids, execute the same binary, and, therefore, end up in the same class using the standard classification criteria.

For more details about application tags, refer to Workload Manager Application Programming Interface.

When an application sets its tag using **wlm\_set\_tag**, it is automatically reclassified according to the current assignment rules and the new tag is taken into account when doing this reclassification.

In addition to the tag itself, the application can also specify flags indicating to WLM if a child process should inherit the tag from its parent after a **fork** or an **exec** subroutine.

A process does not require any special privileges to set its tag.

### Parameters

*tag*

Specifies the address of a character string. An error is returned if this tag is too long.

*flags*

Specifies the address of an integer interpreted in a manner similar to the **versflags** field of the **wlmargs** structure passed to other API routines. The integer pointed to by *flags* should be initialized with **WLM\_VERSION**. In addition, one or more of the following values can be or'ed to **WLM\_VERSION**:

**SWLMTAGINHERITFORK**

Specifies that the children of this process inherit the parent's tag on the **fork** subroutine.

**SWLMTAGINHERITEXEC**

Specifies that the process retains its tag after a call to the **exec** subroutine.

Both flags can be set to specify that the children of a tagged process inherits the tag on the **fork** subroutine and then retains it on the **exec** subroutine.

## Return Values

Upon successful completion, the **wlm\_set\_tag** subroutine returns a value of 0. In case of error, a non-0 value is returned.

## Error Codes

For a list of the possible error codes returned by the WLM API functions, see the description of the **wlm.h** header file.

## Related Information

The **wlm.h** header file.

Workload Manager rules File in *AIX 5L Version 5.2 Files Reference*.

---

## wmemchr Subroutine

### Purpose

Find a wide-character in memory.

### Library

Standard library (**libc.a**)

### Syntax

```
#include <wchar.h>
wchar_t *wmemchr (const wchar_t * ws, wchar_t wc, size_t n) ;
```

### Description

The **wmemchr** function locates the first occurrence of **wc** in the initial **n** wide-characters of the object pointed to be **ws**. This function is not affected by locale and all **wchar\_t** values are treated identically. The null wide-character and **wchar\_t** values not corresponding to valid characters are not treated specially.

If **n** is zero, **ws** must be a valid pointer and the function behaves as if no valid occurrence of **wc** is found.

## Return Values

The **wmemchr** function returns a pointer to the located wide-character, or a null pointer if the wide-character does not occur in the object.

## Related Information

The **wmemcmp** (“wmemcmp Subroutine”) subroutine, **wmemcpy** (“wmemcpy Subroutine”) subroutine, **wmemmove** (“wmemmove Subroutine” on page 502) subroutine, **wmemset** (“wmemset Subroutine” on page 503) subroutine.

---

## wmemcmp Subroutine

### Purpose

Compare wide-characters in memory.

### Library

Standard library (**libc.a**)

### Syntax

```
#include <wchar.h>
int wmemcmp (const wchar_t * ws1, const wchar_t * ws2, size_t n);
```

### Description

The **wmemcmp** function compares the first **n** wide-characters of the object pointed to by **ws1** to the first **n** wide-characters of the object pointed to by **ws2**. This function is not affected by locale and all **wchar\_t** values are treated identically. The null wide-character and **wchar\_t** values not corresponding to valid characters are not treated specially.

If **n** is zero, **ws1** and **ws2** must be a valid pointers and the function behaves as if the two objects compare equal.

### Return Values

The **wmemcmp** function returns an integer greater than, equal to, or less than zero, accordingly as the object pointed to by **ws1** is greater than, equal to, or less than the object pointed to by **ws2**.

## Related Information

The **wmemchr** (“wmemchr Subroutine” on page 500) subroutine, **wmemcpy** (“wmemcpy Subroutine”) subroutine, **wmemmove** (“wmemmove Subroutine” on page 502) subroutine, **wmemset** (“wmemset Subroutine” on page 503) subroutine.

---

## wmemcpy Subroutine

### Purpose

Copy wide-characters in memory.

### Library

Standard library (**libc.a**)

## Syntax

```
#include <wchar.h>
wchar_t *wmemcpy (wchar_t * ws1, const wchar_t * ws2, size_t n) ;
```

## Description

The **wmemcpy** function copies **n** wide-characters from the object pointed to by **ws2** to the object pointed to be **ws1**. This function is not affected by locale and all **wchar\_t** values are treated identically. The null wide-character and **wchar\_t** values not corresponding to valid characters are not treated specially.

If **n** is zero, **ws1** and **ws2** must be a valid pointers, and the function copies zero wide-characters.

## Return Values

The **wmemcpy** function returns the value of **ws1**.

## Related Information

The **wmemchr** (“wmemchr Subroutine” on page 500) subroutine, **wmemcmp** (“wmemcmp Subroutine” on page 501) subroutine, **wmemmove** (“wmemmove Subroutine”) subroutine, **wmemset** (“wmemset Subroutine” on page 503) subroutine.

---

## wmemmove Subroutine

### Purpose

Copy wide-characters in memory with overlapping areas.

### Library

Standard library (**libc.a**)

### Syntax

```
#include <wchar.h>
wchar_t *wmemmove (wchar_t * ws1, const wchar_t * ws2, size_t n) ;
```

### Description

The **wmemmove** function copies **n** wide-characters from the object pointed to by **ws2** to the object pointed to by **ws1**. Copying takes place as if the **n** wide-characters from the object pointed to by **ws2** are first copied into a temporary array of **n** wide-characters that does not overlap the objects pointed to by **ws1** or **ws2**, and then the **n** wide-characters from the temporary array are copied into the object pointed to by **ws1**.

This function is not affected by locale and all **wchar\_t** values are treated identically. The null wide-character and **wchar\_t** values not corresponding to valid characters are not treated specially.

If **n** is zero, **ws1** and **ws2** must be a valid pointers, and the function copies zero wide-characters.

### Return Values

The **wmemmove** function returns the value of **ws1**.

### Related Information

The **wmemchr** (“wmemchr Subroutine” on page 500) subroutine, **wmemcmp** (“wmemcmp Subroutine” on page 501) subroutine, **wmemcpy** (“wmemcpy Subroutine” on page 501) subroutine, **wmemset** (“wmemset Subroutine” on page 503) subroutine.

---

## wmemset Subroutine

### Purpose

Set wide-characters in memory.

### Library

Standard library (**libc.a**)

### Syntax

```
#include <wchar.h>
wchar_t *wmemset (wchar_t * ws, wchar_t wc, size_t n);
```

### Description

The **wmemset** function copies the value of *wc* into each of the first *n* wide-characters of the object pointed to by *ws*. This function is not affected by locale and all *wchar\_t* values are treated identically. The null wide-character and *wchar\_t* values not corresponding to valid characters are not treated specially. If *n* is zero, *ws* must be a valid pointer and the function copies zero wide-characters.

### Return Values

The **wmemset** functions returns the value of *ws*.

### Related Information

The **wmemchr** (“wmemchr Subroutine” on page 500) subroutine, **wmemcmp** (“wmemcmp Subroutine” on page 501) subroutine, **wmemcpy** (“wmemcpy Subroutine” on page 501) subroutine, **wmemmove** (“wmemmove Subroutine” on page 502) subroutine.

---

## wordexp Subroutine

### Purpose

Expands tokens from a stream of words.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <wordexp.h>
int wordexp ( Words, Pwordexp, Flags)
const char *Words;
wordexp_t *Pwordexp;
int Flags;
```

### Description

The **wordexp** subroutine performs word expansions equivalent to the word expansion that would be performed by the shell if the contents of the *Words* parameter were arguments on the command line. The list of expanded words are placed in the *Pwordexp* parameter. The expansions are the same as that which would be performed by the shell if the *Words* parameter were the part of a command line representing the parameters to a command. Therefore, the *Words* parameter cannot contain an unquoted <newline> character or any of the unquoted shell special characters | (pipe), & (ampersand), ; (semicolon), < (less

than sign), or > (greater than sign), except in the case of command substitution. The *Words* parameter also cannot contain unquoted parentheses or braces, except in the case of command or variable substitution. If the *Words* parameter contains an unquoted comment character # (number sign) that is the beginning of a token, the **wordexp** subroutine may treat the comment character as a regular character, or may interpret it as a comment indicator and ignore the remainder of the expression in the *Words* parameter.

The **wordexp** subroutine allows an application to perform all of the shell's expansions on a word or words obtained from a user. For example, if the application prompts for a file name (or a list of file names) and then uses the **wordexp** subroutine to process the input, the user could respond with anything that would be valid as input to the shell.

The **wordexp** subroutine stores the number of generated words and a pointer to a list of pointers to words in the *Pwordexp* parameter. Each individual field created during the field splitting or path name expansion is a separate word in the list specified by the *Pwordexp* parameter. The first pointer after the last last token in the list is a null pointer. The expansion of special parameters \* (asterisk), @ (at sign), # (number sign), ? (question mark), - (minus sign), \$ (dollar sign), ! (exclamation point), and 0 is unspecified.

The words are expanded in the order shown below:

1. Tilde expansion is performed first.
2. Parameter expansion, command substitution, and arithmetic expansion are performed next, from beginning to end.
3. Field splitting is then performed on fields generated by step 2, unless the IFS (input field separators) is full.
4. Path-name expansion is performed, unless the **set -f** command is in effect.
5. Quote removal is always performed last.

## Parameters

<i>Flags</i>	Contains a bit flag specifying the configurable aspects of the <b>wordexp</b> subroutine.
<i>Pwordexp</i>	Contains a pointer to a <b>wordexp_t</b> structure.
<i>Words</i>	Specifies the string containing the tokens to be expanded.

The value of the *Flags* parameter is the bitwise, inclusive OR of the constants below, which are defined in the **wordexp.h** file.

<b>WRDE_APPEND</b>	Appends words generated to those generated by a previous call to the <b>wordexp</b> subroutine.
<b>WRDE_DOOFFS</b>	Makes use of the <b>we_offs</b> structure. If the <b>WRDE_DOOFFS</b> flag is set, the <b>we_offs</b> structure is used to specify the number of null pointers to add to the beginning of the <b>we_words</b> structure. If the <b>WRDE_DOOFFS</b> flag is not set in the first call to the <b>wordexp</b> subroutine with the <i>Pwordexp</i> parameter, it should not be set in subsequent calls to the <b>wordexp</b> subroutine with the <i>Pwordexp</i> parameter.
<b>WRDE_NOCMD</b>	Fails if command substitution is requested.
<b>WRDE_REUSE</b>	The <i>Pwordexp</i> parameter was passed to a previous successful call to the <b>wordexp</b> subroutine. Therefore, the memory previously allocated may be reused.
<b>WRDE_SHOWERR</b>	Does not redirect standard error to <b>/dev/null</b> .
<b>WRDE_UNDEF</b>	Reports 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 the **wordexp** subroutine. The following rules apply when two or more calls to the **wordexp** subroutine are made with the same value of the *Pwordexp* parameter and without intervening calls to the **wordfree** subroutine:

1. The first such call does not set the **WRDE\_APPEND** flag. All subsequent calls set it.

2. For a single invocation of the **wordexp** subroutine, all calls either set the **WRDE\_DOOFFS** flag, or do not set it.
3. After the second and each subsequent call, the *Pwordexp* parameter points to a list containing the following:
  - a. Zero or more null characters, as specified by the **WRDE\_DOOFFS** flag and the **we\_offs** structure.
  - b. Pointers to the words that were in the *Pwordexp* parameter 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 the *Pwordexp* parameter is the total number of words from all of the calls.
5. The application should not modify the *Pwordexp* parameter between the calls.

The **WRDE\_NOCMD** flag is provided for applications that, for security or other reasons, want to prevent a user from executing shell commands. Disallowing unquoted shell special characters also prevents unwanted side effects such as executing a command or writing to a file.

Unless the **WRDE\_SHOWERR** flag is set in the *Flags* parameter, the **wordexp** subroutine redirects standard error to the */dev/null* file for any utilities executed as a result of command substitution while expanding the *Words* parameter. If the **WRDE\_SHOWERR** flag is set, the **wordexp** subroutine may write messages to standard error if syntax errors are detected while expanding the *Words* parameter.

The *Pwordexp* structure is allocated by the caller, but memory to contain the expanded tokens is allocated by the **wordexp** subroutine and added to the structure as needed.

The *Words* parameter cannot contain any `<newline>` characters, or any of the unquoted shell special characters `|`, `&`, `;`, `()`, `{}`, `<`, or `>`, except in the context of command substitution.

## Return Values

If no errors are encountered while expanding the *Words* parameter, the **wordexp** subroutine returns a value of 0. If an error occurs, it returns a nonzero value indicating the error.

## Errors

If the **wordexp** subroutine terminates due to an error, it returns one of the nonzero constants below, which are defined in the **wordexp.h** file.

<b>WRDE_BADCHAR</b>	One of the unquoted characters <code> </code> , <code>&amp;</code> , <code>;</code> , <code>&lt;</code> , <code>&gt;</code> , parenthesis, or braces appears in the <i>Words</i> parameter in an inappropriate context.
<b>WRDE_BADVAL</b>	Reference to undefined shell variable when the <b>WRDE_UNDEF</b> flag is set in the <i>Flags</i> parameter.
<b>WRDE_CMDSUB</b>	Command substitution requested when the <b>WRDE_NOCMD</b> flag is set in the <i>Flags</i> parameter.
<b>WRDE_NOSPACE</b>	Attempt to allocate memory was unsuccessful.
<b>WRDE_SYNTAX</b>	Shell syntax error, such as unbalanced parentheses or unterminated string.

If the **wordexp** subroutine returns the error value **WRDE\_SPACE**, then the expression in the *Pwordexp* parameter is updated to reflect any words that were successfully expanded. In other cases, the *Pwordexp* parameter is not modified.

## Related Information

The **glob** subroutine, **wordfree** (“wordfree Subroutine” on page 506) subroutine.

For more information on basic and extended regular expressions, see *Manipulating Strings with sed*.

---

## wordfree Subroutine

### Purpose

Frees all memory associated with the *Pwordexp* parameter.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <wordexp.h>
```

```
void wordfree ( Pwordexp)  
wordexp_t *Pwordexp;
```

### Description

The **wordfree** subroutine frees any memory associated with the *Pwordexp* parameter from a previous call to the **wordexp** subroutine.

### Parameters

*Pwordexp*        Structure containing a list of expanded words.

### Related Information

The **wordexp** (“wordexp Subroutine” on page 503) subroutine.

---

## write, writex, writev, writevx or pwrite Subroutines

### Purpose

Writes to a file.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
ssize_t write (FileDescriptor, Buffer, NBytes)  
int FileDescriptor;  
const void * Buffer;  
size_t NBytes;
```

```
int writex (FileDescriptor, Buffer, NBytes, Extension)  
int FileDescriptor;  
char *Buffer;  
unsigned int NBytes;  
int Extension;  
#include <sys/uio.h>
```

```

ssize_t write (FileDescriptor, iov, iovCount)
int FileDescriptor;
const struct iovec * iov;
int iovCount;

ssize_t writev (FileDescriptor, iov, iovCount, Extension)
int FileDescriptor;
struct iovec *iov;
int iovCount;
int Extension;

```

```

ssize_t pwrite (FileDescriptor, Buffer, NBytes, Offset)
int FileDescriptor;
const void * Buffer;
size_t NBytes;
off_t Offset;

```

## Description

The **write** subroutine attempts to write the number of bytes of data specified by the *NBytes* parameter to the file associated with the *FileDescriptor* parameter from the buffer pointed to by the *Buffer* parameter.

The **writev** subroutine performs the same action but gathers the output data from the *iovCount* buffers specified by the array of **iovec** structures pointed to by the *iov* parameter. Each **iovec** entry specifies the base address and length of an area in memory from which data should be written. The **writev** subroutine always writes a complete area before proceeding to the next.

The **writex** and **writevx** subroutines are the same as the **write** and **writev** subroutines, respectively, with the addition of an *Extension* parameter, which is used when writing to some device drivers.

With regular files and devices capable of seeking, the actual writing of data proceeds from the position in the file indicated by the file pointer. Upon return from the **write** subroutine, the file pointer increments by the number of bytes actually written.

With devices incapable of seeking, writing always takes place starting at the current position. The value of a file pointer associated with such a device is undefined.

If a **write** requests that more bytes be written than there is room for (for example, the **ulimit** or the physical end of a medium), only as many bytes as there is room for will be written. For example, suppose there is space for 20 bytes more in a file before reaching a limit. A write of 512 bytes will return 20. The next write of a non-zero number of bytes will give a failure return (except as noted below) and the implementation will generate a SIGXFSZ signal for the thread.

Fewer bytes can be written than requested if there is not enough room to satisfy the request. In this case the number of bytes written is returned. The next attempt to write a nonzero number of bytes is unsuccessful (except as noted in the following text). The limit reached can be either that set by the **ulimit** subroutine or the end of the physical medium.

Successful completion of a **write** subroutine clears the SetUserID bit (**S\_ISUID**) of a file if all of the following are true:

- The calling process does not have root user authority.
- The effective user ID of the calling process does not match the user ID of the file.
- The file is executable by the group (**S\_IXGRP**) or other (**S\_IXOTH**).

The **write** subroutine clears the SetGroupID bit (**S\_ISGID**) if all of the following are true:

- The calling process does not have root user authority.

- The group ID of the file does not match the effective group ID or one of the supplementary group IDs of the process.
- The file is executable by the owner (**S\_IXUSR**) or others (**S\_IXOTH**).

**Note:** Clearing of the SetUserID and SetGroupID bits can occur even if the **write** subroutine is unsuccessful, if file data was modified before the error was detected.

If the **O\_APPEND** flag of the file status is set, the file offset is set to the end of the file prior to each write.

If the *FileDescriptor* parameter refers to a regular file whose file status flags specify **O\_SYNC**, this is a synchronous update (as described in the **open** subroutine).

If the *FileDescriptor* parameter refers to a regular file that a process has opened with the **O\_DEFER** file status flag set, the data and file size are not updated on permanent storage until a process issues an **fsync** subroutine or performs a synchronous update. If all processes that have the file open with the **O\_DEFER** file status flag set close the file before a process issues an **fsync** subroutine or performs a synchronous update, the data and file size are not updated on permanent storage.

Write requests to a pipe (or first-in-first-out (FIFO)) are handled the same as a regular file with the following exceptions:

- There is no file offset associated with a pipe; hence, each write request appends to the end of the pipe.
- If the size of the write request is less than or equal to the value of the **PIPE\_BUF** system variable (described in the **pathconf** routine), the **write** subroutine is guaranteed to be atomic. The data is not interleaved with data from other write processes on the same pipe. Writes of greater than **PIPE\_BUF** bytes can have data interleaved, on arbitrary boundaries, with writes by other processes, whether or not the **O\_NDELAY** or **O\_NONBLOCK** file status flags are set.
- If the **O\_NDELAY** and **O\_NONBLOCK** file status flags are clear (the default), a write request to a full pipe causes the process to block until enough space becomes available to handle the entire request.
- If the **O\_NDELAY** file status flag is set, a write to a full pipe returns a 0.
- If the **O\_NONBLOCK** file status flag is set, a write to a full pipe returns a value of -1 and sets the **errno** global variable to **EAGAIN**.

When attempting to write to a character special file that supports nonblocking writes and no data can currently be written (streams are an exception described later in this article):

- If the **O\_NDELAY** and **O\_NONBLOCK** flags are clear (the default), the **write** subroutine blocks until data can be written.
- If the **O\_NDELAY** flag is set, the **write** subroutine returns 0.
- If the **O\_NONBLOCK** flag is set, the **write** subroutine returns -1 and sets the **errno** global variable to **EAGAIN** if no data can be written.

When attempting to write to a regular file that supports enforcement-mode record locks, and all or part of the region to be written is currently locked by another process, the following can occur:

- If the **O\_NDELAY** and **O\_NONBLOCK** file status flags are clear (the default), the calling process blocks until the lock is released.
- If the **O\_NDELAY** or **O\_NONBLOCK** file status flag is set, then the **write** subroutine returns a value of -1 and sets the **errno** global variable to **EAGAIN**.

**Note:** The **fcntl** subroutine provides more information about record locks.

If **fildes** refers to a STREAM, the operation of **write** is determined by the values of the minimum and maximum nbyte range ("packet size") accepted by the STREAM. These values are determined by the topmost STREAM module. If nbyte falls within the packet size range, nbyte bytes will be written. If nbyte does not fall within the range and the minimum packet size value is 0, **write** will break the buffer into

maximum packet size segments prior to sending the data downstream (the last segment may contain less than the maximum packet size). If `nbyte` does not fall within the range and the minimum value is non-zero, **write** will fail with `errno` set to **ERANGE**. Writing a zero-length buffer (`nbyte` is 0) to a STREAMS device sends 0 bytes with 0 returned. However, writing a zero-length buffer to a STREAMS-based pipe or FIFO sends no message and 0 is returned. The process may issue `I_SWROPT ioctl` to enable zero-length messages to be sent across the pipe or FIFO.

When writing to a STREAM, data messages are created with a priority band of 0. When writing to a STREAM that is not a pipe or FIFO:

- **O\_NONBLOCK** should specify either **O\_NONBLOCK** or **O\_NDELAY**. The IBM streams implementation treats these two the same.
- If **O\_NONBLOCK** or **O\_NDELAY** is clear, and the STREAM cannot accept data (the STREAM write queue is full due to internal flow control conditions), **write** will block until data can be accepted.
- If **O\_NONBLOCK** or **O\_NDELAY** is set and the STREAM cannot accept data, **write** will return **-1** and set `errno` to **EAGAIN**.
- If **O\_NONBLOCK** or **O\_NDELAY** is set and part of the buffer has been written while a condition in which the STREAM cannot accept additional data occurs, **write** will terminate and return the number of bytes written.

**Note:** The IBM streams implementation treats **O\_NONBLOCK** and **O\_NDELAY** the same.

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

The **writew** function is equivalent to **write**, but gathers the output data from the `iovcnt` buffers specified by the members of the `iov` array: `iov[0]`, `iov[1]`, ..., `iov[iovcnt - 1]`. **iovcnt** is valid if greater than 0 and less than or equal to `{IOV_MAX}`, defined in **limits.h**.

Each **iovec** entry specifies the base address and length of an area in memory from which data should be written. The **writew** function will always write a complete area before proceeding to the next.

If **files** refers to a regular file and all of the **iov\_len** members in the array pointed to by **iov** are 0, **writew** will return 0 and have no other effect. For other file types, the behaviour is unspecified.

If the sum of the **iov\_len** values is greater than `SSIZE_MAX`, the operation fails and no data is transferred.

The behavior of an interrupted **write** subroutine depends on how the handler for the arriving signal was installed. The handler can be installed in one of two ways, with the following results:

- If the handler was installed with an indication that subroutines should not be restarted, the **write** subroutine returns a value of **-1** and sets the `errno` global variable to **EINTR** (even if some data was already written).
- If the handler was installed with an indication that subroutines should be restarted, and:
  - If no data had been written when the interrupt was handled, the **write** subroutine will not return a value (it is restarted).
  - If data had been written when the interrupt was handled, this **write** subroutine returns the amount of data already written.

**Note:** A write to a regular file is not interruptible. Only writes to objects that may block indefinitely, such as FIFOs, sockets, and some devices, are generally interruptible.

The **pwrite** function performs the same action as **write**, except that it writes into a given position without changing the file pointer. The first three arguments to **pwrite** are the same as **write** with the addition of a fourth argument offset for the desired position inside the file.

**Note:** The **pwrite64** subroutine applies to AIX 4.3 and later.

```
ssize_t pwrite64(int fd , const void *buf , size_t nbytes , off64_t offset)
```

The **pwrite64** subroutine performs the same action as **pwrite** but the limit of offset to the maximum file size for the file associated with the `fileDescriptor` and `DEV_OFF_MAX` if the file associated with `fileDescriptor` is a block special or character special file.

## Parameters

<i>Buffer</i>	Identifies the buffer containing the data to be written.
<i>Extension</i>	Provides communication with character device drivers that require additional information or return additional status. Each driver interprets the <i>Extension</i> parameter in a device-dependent way, either as a value or as a pointer to a communication area. Drivers must apply reasonable defaults when the <i>Extension</i> parameter value is 0.
<i>FileDescriptor</i>	Identifies the object to which the data is to be written.
<i>iov</i>	Points to an array of <b>iovec</b> structures, which identifies the buffers containing the data to be written. The <b>iovec</b> structure is defined in the <b>sys/uio.h</b> file and contains the following members: <pre>caddr_t iov_base; size_t iov_len;</pre>
<i>iovCount</i>	Specifies the number of <b>iovec</b> structures pointed to by the <i>iov</i> parameter.
<i>NBytes</i>	Specifies the number of bytes to write.

## Return Values

Upon successful completion, the **write**, **writex**, **writev**, and **writevx** subroutines return the number of bytes that were actually written. The number of bytes written is never greater than the value specified by the *NBytes* parameter. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **write**, **writex**, **writev**, and **writevx** subroutines are unsuccessful when one of the following is true:

<b>EAGAIN</b>	The <b>O_NONBLOCK</b> flag is set on this file and the process would be delayed in the write operation; or an enforcement-mode record lock is outstanding in the portion of the file that is to be written.
<b>EBADF</b>	The <i>FileDescriptor</i> parameter does not specify a valid file descriptor open for writing.
<b>EDQUOT</b>	New disk blocks cannot be allocated for the file because the user or group quota of disk blocks has been exhausted on the file system.
<b>EFBIG</b>	An offset greater than <b>MAX_FILESIZE</b> was requested on the 32-bit kernel.
<b>EFAULT</b>	The <i>Buffer</i> parameter or part of the <i>iov</i> parameter points to a location outside of the allocated address space of the process.
<b>EFBIG</b>	An attempt was made to write a file that exceeds the process' file size limit or the maximum file size. If the user has set the environment variable <b>XPG_SUS_ENV=ON</b> prior to execution of the process, then the <b>SIGXFSZ</b> signal is posted to the process when exceeding the process' file size limit.
<b>EINVAL</b>	The file position pointer associated with the <i>FileDescriptor</i> parameter was negative; the <i>iovCount</i> parameter value was not between 1 and 16, inclusive; or one of the <b>iov_len</b> values in the <b>iov</b> array was negative or the sum overflowed a 32-bit integer.
<b>EINVAL</b>	The STREAM or multiplexer referenced by <i>FileDescriptor</i> is linked (directly or indirectly) downstream from a multiplexer.
<b>EINVAL</b>	The value of the <i>Nbytes</i> parameter that is larger than <b>OFF_MAX</b> , was requested on the 32-bit kernel. This is a case where the system call is requested from a 64-bit application that is running on a 32-bit kernel.
<b>EINTR</b>	A signal was caught during the write operation, and the signal handler was installed with an indication that subroutines are not to be restarted.

<b>EIO</b>	An I/O error occurred while writing to the file system; or the process is a member of a background process group attempting to write to its controlling terminal, <b>TOSTOP</b> is set, the process is neither ignoring nor blocking <b>SIGTTOU</b> , and the process group has no parent process.
<b>ENOSPC</b>	No free space is left on the file system containing the file.
<b>ENXIO</b>	A hangup occurred on the STREAM being written to.
<b>EPIPE</b>	An attempt was made to write to a file that is not opened for reading by any process, or to a socket of type <b>SOCK_STREAM</b> that is not connected to a peer socket; or an attempt was made to write to a pipe or FIFO that is not open for reading by any process. If this occurs, a <b>SIGPIPE</b> signal will also be sent to the process.
<b>ERANGE</b>	The transfer request size was outside the range supported by the STREAMS file associated with <i>FileDescriptor</i> .

The **write**, **writex**, **writexv**, and **writexv** subroutines may be unsuccessful if the following is true:

<b>ENXIO</b>	A request was made of a nonexistent device, or the request was outside the capabilities of the device.
<b>EFBIG</b>	An attempt was made to write to a regular file where NBytes greater than zero and the starting offset is greater than or equal to the offset maximum established in the open file description associated with <i>FileDescriptor</i> .
<b>EINVAL</b>	The <b>offset</b> argument is invalid. The value is negative.
<b>ESPIPE</b>	<b>fildes</b> is associated with a pipe or FIFO.

## Related Information

The **fcntl**, **dup**, or **dup2** subroutine, **fsync** subroutine, **ioctl** subroutine, **lockfx** subroutine, **lseek** subroutine, **open**, **openx**, or **creat** subroutine, **pathconf** subroutine, **pipe** subroutine, **poll** subroutine, **select** (“select Subroutine” on page 115) subroutine, **ulimit** (“ulimit Subroutine” on page 415) subroutine.

The **limits.h** file, **unistd.h** file.

The Input and Output Handling Programmer’s Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## wstring Subroutine

### Purpose

Perform operations on wide character strings.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <wstring.h>
```

```
wchar_t *wstrcat ("wstring Subroutine") (XString1, XString2)
wchar_t *XString1, *XString2;
```

```
wchar_t * wstrncat (XString, XString2, Number)
wchar_t *XString1, *XString2;
int Number;
```

```
int wstrcmp (XString1, XString2)
wchar_t *XString1, *XString2;
```

```

int wstrncmp (XString1, XString2, Number)
wchar_t *XString1, *XString2;
int Number;

wchar_t * wstrncpy (XString1, XString2)
wchar_t *XString1, *XString2;

wchar_t * wstrncpy (XString1, XString2, Number)
wchar_t *XString1, *XString2;
int Number;

int wstrlen (XString)
wchar_t *XString;

wchar_t * wstrchr (XString, Number)
wchar_t *XString;
int Number;

wchar_t * wstrrchr (XString, Number)
wchar_t *XString;
int Number;

wchar_t * wstrpbrk (XString1, XString2)
wchar_t *XString1, XString2;

int wstrspn (XString1, XString2)
wchar_t *XString1, XString2;

int wstrcspn (XString1, XString2)
wchar_t *XString1, XString2;

wchar_t * wstrtok (XString1, XString2)
wchar_t *XString1, XString2;

wchar_t * wstrdup (XString1)
wchar_t *XString1;

```

## Description

The **wstring** subroutines copy, compare, and append strings in memory, and determine location, size, and existence of strings in memory. For these subroutines, a string is an array of **wchar\_t** characters, terminated by a null character. The **wstring** subroutines parallel the **string** subroutines, but operate on strings of type **wchar\_t** rather than on type **char**, except as specifically noted below.

The parameters *XString1*, *XString2*, and *XString* point to strings of type **wchar\_t** (arrays of **wchar** characters terminated by a **wchar\_t** null character).

The subroutines **wstrcat**, **wstrncat**, **wstrncpy**, and **wstrncpy** all alter the *XString1* parameter. They do not check for overflow of the array pointed to by *XString1*. All string movement is performed wide character by wide character. Overlapping moves toward the left work as expected, but overlapping moves to the right may give unexpected results. All of these subroutines are declared in the **wstring.h** file.

The **wstrcat** subroutine appends a copy of the **wchar\_t** string pointed to by the *XString2* parameter to the end of the **wchar\_t** string pointed to by the *XString1* parameter. The **wstrcat** subroutine returns a pointer to the null-terminated result.

The **wstrncat** subroutine copies, at most, the value of the *Number* parameter of **wchar\_t** characters in the *XString2* parameter to the end of the **wchar\_t** string pointed to by the *XString1* parameter. Copying stops before *Number* **wchar\_t** character if a null character is encountered in the string pointed to by the *XString2* parameter. The **wstrncat** subroutine returns a pointer to the null-terminated result.

The **wstricmp** subroutine lexicographically compares the **wchar\_t** string pointed to by the *XString1* parameter to the **wchar\_t** string pointed to by the *XString2* parameter. The **wstricmp** subroutine returns a value that is:

- Less than 0 if *XString1* is less than *XString2*
- Equal to 0 if *XString1* is equal to *XString2*
- Greater than 0 if *XString1* is greater than *XString2*

The **wstrncmp** subroutine makes the same comparison as **wstricmp**, but it compares, at most, the value of the *Number* parameter of pairs of **wchar** characters. The comparisons are based on collation values as determined by the locale category **LC\_COLLATE** and the **LANG** variable.

The **wstrcpy** subroutine copies the string pointed to by the *XString2* parameter to the array pointed to by the *XString1* parameter. Copying stops when the **wchar\_t** null is copied. The **wstrcpy** subroutine returns the value of the *XString1* parameter.

The **wstrncpy** subroutine copies the value of the *Number* parameter of **wchar\_t** characters from the string pointed to by the *XString2* parameter to the **wchar\_t** array pointed to by the *XString1* parameter. If *XString2* is less than *Number* **wchar\_t** characters long, then **wstrncpy** pads *XString1* with trailing null characters to fill *Number* **wchar\_t** characters. If *XString2* is *Number* or more **wchar\_t** characters long, only the first *Number* **wchar\_t** characters are copied; the result is not terminated with a null character. The **wstrncpy** subroutine returns the value of the *XString1* parameter.

The **wstrlen** subroutine returns the number of **wchar\_t** characters in the string pointed to by the *XString* parameter, not including the terminating **wchar\_t** null.

The **wstrchr** subroutine returns a pointer to the first occurrence of the **wchar\_t** specified by the *Number* parameter in the **wchar\_t** string pointed to by the *XString* parameter. A null pointer is returned if the **wchar\_t** does not occur in the **wchar\_t** string. The **wchar\_t** null that terminates a string is considered to be part of the **wchar\_t** string.

The **wstrrchr** subroutine returns a pointer to the last occurrence of the character specified by the *Number* parameter in the **wchar\_t** string pointed to by the *XString* parameter. A null pointer is returned if the **wchar\_t** does not occur in the **wchar\_t** string. The **wchar\_t** null that terminates a string is considered to be part of the **wchar\_t** string.

The **wstrpbrk** subroutine returns a pointer to the first occurrence in the **wchar\_t** string pointed to by the *XString1* parameter of any code point from the string pointed to by the *XString2* parameter. A null pointer is returned if no character matches.

The **wstrspn** subroutine returns the length of the initial segment of the string pointed to by the *XString1* parameter that consists entirely of code points from the **wchar\_t** string pointed to by the *XString2* parameter.

The **wstrcspn** subroutine returns the length of the initial segment of the **wchar\_t** string pointed to by the *XString1* parameter that consists entirely of code points *not* from the **wchar\_t** string pointed to by the *XString2* parameter.

The **wstrtok** subroutine returns a pointer to an occurrence of a text token in the string pointed to by the *XString1* parameter. The *XString2* parameter specifies a set of code points as token delimiters. If the *XString1* parameter is anything other than null, then the **wstrtok** subroutine reads the string pointed to by

the *XString1* parameter until it finds one of the delimiter code points specified by the *XString2* parameter. It then stores a **wchar\_t** null into the **wchar\_t** string, replacing the delimiter code point, and returns a pointer to the first **wchar\_t** of the text token. The **wstrtok** subroutine keeps track of its position in the **wchar\_t** string so that subsequent calls with a null *XString1* parameter step through the **wchar\_t** string. The delimiters specified by the *XString2* parameter can be changed for subsequent calls to **wstrtok**. When no tokens remain in the **wchar\_t** string pointed to by the *XString1* parameter, the **wstrtok** subroutine returns a null pointer.

The **wstrdup** subroutine returns a pointer to a **wchar\_t** string that is a duplicate of the **wchar\_t** string to which the *XString1* parameter points. Space for the new string is allocated using the **malloc** subroutine. When a new string cannot be created, a null pointer is returned.

## Related Information

The **malloc** subroutine, **strcat**, **strncat**, **strxfrm**, **strcpy**, **strncpy**, or **strdup** (“strcat, strncat, strxfrm, strcpy, strncpy, or strdup Subroutine” on page 281) subroutine, **strcmp**, **strncmp**, **strcasemp**, **strncasemp**, or **strcoll** (“strcmp, strncmp, strcasemp, strncasemp, or strcoll Subroutine” on page 283) subroutine, **strlen**, **strchr**, **strchr**, **strpbrk**, **strspn**, **strcspn**, **strstr**, or **strtok** (“strlen, strchr, strchr, strpbrk, strspn, strcspn, strstr, strtok, or strsep Subroutine” on page 291) subroutine.

List of String Manipulation Services in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

National Language Support Overview in *AIX 5L Version 5.2 National Language Support Guide and Reference*.

Subroutines, Example Programs, and Libraries in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## wstrtod or watof Subroutine

### Purpose

Converts a string to a double-precision floating-point.

### Library

Standard C Library

### Syntax

```
#include <wstring.h>
```

```
double wstrtod ( String, Pointer)
```

```
wchar_t *String, **Pointer;
```

```
double watof (String)
```

```
wchar_t *String;
```

### Description

The **wstrtod** subroutine returns a double-precision floating-point number that is converted from an **wchar\_t** string pointed to by the *String* parameter. The system searches the *String* until it finds the first unrecognized character.

The **wstrtod** subroutine recognizes a string that starts with any number of white-space characters (defined by the **iswspace** subroutine), followed by an optional sign, a string of decimal digits that may include a decimal point, e or E, an optional sign or space, and an integer.

When the value of *Pointer* is not (**wchar\_t** \*\*) null, a pointer to the search terminating character is returned to the address indicated by *Pointer*. When the resulting number cannot be created, *\*Pointer* is set to *String* and 0 (zero) is returned.

The **watof** (*String*) subroutine functions like the **wstrtod** (*String* (**wchar\_t** \*\*) null).

## Parameters

*String*            Specifies the address of the string to scan.  
*Pointer*           Specifies the address at which the pointer to the terminating character is stored.

## Error Codes

When the value causes overflow, **HUGE\_VAL** (defined in the **math.h** file) is returned with the appropriate sign, and the **errno** global variable is set to **ERANGE**. When the value causes underflow, 0 is returned and the **errno** global variable is set to **ERANGE**.

## Related Information

The **atof**, **atoff**, **strtod**, **strtof** subroutine, **scanf**, **fscanf**, **sscanf** (“scanf, fscanf, sscanf, or wsscanf Subroutine” on page 109) subroutine, **strtol**, **strtoul**, **atol**, **atoi** (“strtol, strtoul, strtoll, strtoull, or atoi Subroutine” on page 299) subroutine, **wstrtol**, **watol**, **watoi** (“wstrtol, watol, or watoi Subroutine”) subroutine.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## wstrtol, watol, or watoi Subroutine

### Purpose

Converts a string to an integer.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <wstring.h>
```

```
long wstrtol ( String, Pointer, Base)  
wchar_t *String, **Pointer;  
int Base;  
long watol (String)  
wchar_t *String;  
int watoi (String)  
wchar_t *String;
```

### Description

The **wstrtol** subroutine returns a long integer that is converted from the string pointed to by the *String* parameter. The string is searched until a character is found that is inconsistent with *Base*. Leading white-space characters defined by the **ctype** subroutine **iswspace** are ignored.

When the value of *Pointer* is not (**wchar\_t \*\***) null, a pointer to the terminating character is returned to the address indicated by *Pointer*. When an integer cannot be created, the address indicated by *Pointer* is set to *String*, and 0 is returned.

When the value of *Base* is positive and not greater than 36, that value is used as the base during conversion. Leading zeros that follow an optional leading sign are ignored. When the value of *Base* is 16, 0x and 0X are ignored.

When the value of *Base* is 0, the system chooses an appropriate base after examining the actual string. An optional sign followed by a leading zero signifies octal, and a leading 0x or 0X signifies hexadecimal. In all other cases, the subroutines assume a decimal base.

Truncation from **long** data type to **int** data type occurs by assignment, and also by explicit casting.

The **watol** (*String*) subroutine functions like **wstrtol** (*String*, (**wchar\_t \*\***) null, **10**).

The **watoi** (*String*) subroutine functions like **(int) wstrtol** (*String*, (**wchar\_t \*\***) null, **10**).

**Note:** Even if overflow occurs, it is ignored.

## Parameters

<i>String</i>	Specifies the address of the string to scan.
<i>Pointer</i>	Specifies the address at which the pointer to the terminating character is stored.
<i>Base</i>	Specifies an integer value used as the base during conversion.

## Related Information

The **atof**, **atoff**, **strtod**, **strtof** subroutine, **scanf**, **fscanf**, **sscanf** (“scanf, fscanf, sscanf, or wsscanf Subroutine” on page 109) subroutine, **strtol**, **strtoul**, **atol**, **atoi** (“strtol, strtoul, strtoll, strtoull, or atoi Subroutine” on page 299) subroutine, **wstrtod**, **watof** (“wstrtod or watof Subroutine” on page 514) subroutine.

Subroutines Overview in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## **xcrypt\_key\_setup, xcrypt\_encrypt, xcrypt\_decrypt, xcrypt\_hash, xcrypt\_malloc, xcrypt\_free, xcrypt\_printb, xcrypt\_btoa and xcrypt\_randbuff Subroutine**

### Purpose

Provides various block and stream cipher algorithms and two crypto-secure hash algorithms.

### Library

Cryptographic Library (**libmodcrypt.a**)

### Syntax

```
#include <xcrypt.h>
```

```
int xcrypt_key_setup (alg, key, keymat, keysize, dir)
int alg;
```

```

xcrypt_key *key;
u_char *keymat;
int keysize;
int dir;

int xcrypt_encrypt (alg, mode, key, IV, in, insize, out, padding)
int alg;
int mode;
xcrypt_key *key;
u_char *IV;
u_char *in;
int insize;
u_char *out;
int padding;

int xcrypt_decrypt (alg, mode, key, IV, in, insize, out, padding)
int alg;
int mode;
xcrypt_key *key;
u_char *IV;
u_char *in;
int insize;
u_char *out;
int padding;

int xcrypt_hash (alg, in, insize, out)
int alg;
u_char *in;
int insize;
u_char *out;

int xcrypt_malloc (pp, size, blocksize)
uchar **pp;
int size;
int blocksize;

void xcrypt_free (p, size)
void *p;
int size;

void xcrypt_printb (p, size)
void *p;
int size;

void xcrypt_btoa (dest, buff, size)
char *dest;
void *buff;
int size;

void xcrypt_randbuff (dest, size)
void *dest;
int size;

```

## Description

These subroutines provide block and stream cipher algorithms, plus two crypto-secure hash algorithms. Encryption may be done through the Rijndael, Mars, and Twofish block ciphers or the SEAL stream cipher. Each of these algorithms uses a use a block length of 128 bits and key lengths of 128, 192 and 256 bits. SEAL is a stream cipher that uses a 160 bit key and a 32 bit word input stream. In addition, the MD5 and SHA-1 cryptographic hash algorithms are included.

The **xcrypt\_key\_setup** subroutine is used to setup a key schedule for any of the block cipher algorithms. It stores the key schedule in the **xcrypt\_key** data structure that is passed in. Note that when using the Twofish method, the *keymat* parameter should be set to NULL.

The **xcrypt\_encrypt** subroutine encrypts a buffer. Data can be encrypted using the CBC mode (Cipher Block Chaining), EBC mode (Electronic Codebook) or CBF1 mode. Note that when EBC mode is being used, no initialization vector is required.

The **xcrypt\_decrypt** subroutine decrypts a buffer. Data can be encrypted using the CBC mode (Cipher Block Chaining), EBC mode (Electronic Codebook) or CBF1 mode. Note that when EBC mode is being used, no initialization vector is required.

The **xcrypt\_hash** subroutine hashes a buffer using either the MD5 or SHA-1 algorithm.

The **xcrypt\_malloc** subroutine dynamically allocates the least size bytes of memory to provide blocks of *blocksize* bytes. For example, if *size* is 105 and *blocksize* is 10, the **xcrypt\_malloc** subroutine will return at least 110 bytes of memory (11 blocks, each 10 bytes in size). The **xcrypt\_malloc** subroutine should be used when you need **xcrypt** to pad buffers. It will make sure that enough memory is allocated for the data to be encrypted, plus the padding.

The **xcrypt\_free** subroutine overwrites and frees dynamically allocated memory.

The **xcrypt\_printb** subroutine prints a buffer to the screen in hexadecimal notation.

The **xcrypt\_btoa** subroutine returns a string representing the buffer in hexadecimal. Note that the *dest* parameter must point to a buffer of  $size * 2 + 1$ .

The **xcrypt\_randbuff** subroutine fills a buffer with random data.

## Parameters

<i>alg</i>	Specifies the cipher to use. Use the symbolic constants that are described below: <b>RIJNDAEL</b> Rijndael (AES) block cipher <b>MARS</b> Mars block cipher <b>TWOFISH</b> Twofish block cipher <b>SEAL</b> SEAL stream cipher <b>SHA1</b> SHA-1 one-way hash function <b>MD5</b> MD5 one-way hash function
<i>key</i>	Points to the key instance to set up. Use for encryption or decryption.
<i>keymat</i>	Points to the key material used to build the key schedule.
<i>keysize</i>	Size of the <i>keymat</i> parameter. Use the symbolic constants described below: <b>KEY_80</b> 80 bit key <b>KEY_128</b> 128 bit key <b>KEY_192</b> 192 bit key <b>KEY_256</b> 256 bit key
<i>dir</i>	The direction (encryption or decryption). Use the symbolic constants described below: <b>DIR_ENCRYPT</b> Encrypt <b>DIR_DECRYPT</b> Decrypt

<i>mode</i>	Specifies the mode of operation. Use the symbolic constants described below: <b>MODE_ECB</b> Ciphering in ECB mode <b>MODE_CBC</b> Ciphering in CBC mode <b>MODE_CFB1</b> Ciphering in 1-bit CFB mode
<i>IV</i>	Points to the buffer holding the initialization vector. <b>Note:</b> When using ECB mode, the <i>IV</i> parameter should point to NULL.
<i>in</i>	Points to the buffer holding the data to encrypt, decrypt, or hash.
<i>insize</i>	Contains the size of the <i>in</i> parameter.
<i>out</i>	Points to a preallocated output buffer.
<i>padding</i>	Specifies whether <b>xcrypt</b> should pad the buffers or not. Use the symbolic constants described below: <b>TRUE</b> True <b>FALSE</b> False
<i>pp</i>	A double pointer to the destination.
<i>size</i>	Contains the amount of memory to allocate, deallocate, print the contents of, or convert to a string.
<i>blocksize</i>	Contains the size of the blocks. Use the symbolic constants described below: <b>BITS_32</b> 32 bits <b>BITS_80</b> 80 bits <b>BITS_128</b> 128 bits <b>BITS_160</b> 160 bits <b>BITS_192</b> 192 bits <b>BITS_256</b> 256 bits
<i>p</i>	Points to the memory to overwrite and free.
<i>buff</i>	Points to a buffer to print or convert to a string.
<i>dest</i>	Points to a preallocated destination buffer.

## Return Values

The **xcrypt\_key\_setup** and **xcrypt\_hash** subroutines return 0 on success. The **xcrypt\_malloc** subroutine returns the amount of memory allocated on success. The **xcrypt\_encrypt** subroutine returns the amount of data encrypted on success. The **xcrypt\_decrypt** subroutine returns the amount of data decrypted on success.

On failure the above subroutines return the following error codes:

## Error Codes

**xcrypt\_key\_setup:**

<b>BAD_ALIGN32</b>	A parameter is not aligned on a 32 bit boundary.
<b>BAD_KEY_DIR</b>	The <i>dir</i> parameter is not valid
<b>BAD_KEY_INSTANCE</b>	The <i>key</i> parameter is not valid

**BAD\_KEY\_MAT**

The *keysize* parameter is not valid or the *key* parameter is corrupt.

**xcrypt\_encrypt:**

**BAD\_ALG**

The *alg* parameter is not valid.

**BAD\_CIPHER\_MODE**

The *mode* parameter is not valid.

**BAD\_CIPHER\_STATE**

The *key* parameter is not valid.

**BAD\_INPUT\_LEN**

The *insize* parameter is not a multiple of the *blocksize* being used by a block cipher for encryption or decryption.

**BAD\_IV**

The *IV* parameter is set to NULL when the *mode* parameter is set to **MODE\_CBC**.

**BAD\_IV\_MAT**

The *IV* parameter is not valid.

**BAD\_KEY\_INSTANCE**

The *key* parameter is not valid.

**xcrypt\_decrypt:**

**BAD\_ALG**

The *alg* parameter is not valid.

**BAD\_CIPHER\_MODE**

The *mode* parameter is not valid.

**BAD\_CIPHER\_STATE**

The *key* parameter is not valid.

**BAD\_INPUT\_LEN**

The *insize* parameter is not a multiple of the *blocksize* being used by a block cipher for encryption or decryption.

**BAD\_IV**

The *IV* parameter is set to NULL when the *mode* parameter is set to **MODE\_CBC**.

**BAD\_IV\_MAT**

The *IV* parameter is not valid.

**BAD\_KEY\_INSTANCE**

The *key* parameter is not valid.

**xcrypt\_hash:**

**BAD\_ALG**

The *alg* parameter is not valid.

**xcrypt\_malloc:**

**BAD\_MEM\_ALLOC**

The system could not allocate *size* bytes.

---

## yield Subroutine

### Purpose

Yields the processor to processes with higher priorities.

### Library

Standard C library (**libc.a**)

### Syntax

```
void yield (void);
```

### Description

The **yield** subroutine forces the current running process or thread to relinquish use of the processor. If the run queue is empty when the **yield** subroutine is called, the calling process or kernel thread is immediately

rescheduled. If the calling process has multiple threads, only the calling thread is affected. The process or thread resumes execution after all threads of equal or greater priority are scheduled to run.

## **Related Information**

The **getpriority**, **setpriority**, or **nice** subroutine, **setpri** (“setpri Subroutine” on page 147) subroutine.



---

## Chapter 2. Curses Subroutines

---

### addch, mvaddch, mvwaddch, or waddch Subroutine

#### Purpose

Adds a single-byte character and rendition to a window and advances the cursor.

#### Library

Curses Library (**libcurses.a**)

#### Syntax

```
#include <curses.h>
int addch(const chtype ch);
int mvaddch(int y,
int x,
const chtype ch);
int mvwaddch(WINDOW *in,
const chtype ch);
int waddch(WINDOW *win,
const chtype ch);
```

#### Description

The **addch**, **waddch**, **mvaddch**, and **mvwaddch** subroutines add a character to a window at the logical cursor location. After adding the character, curses advances the position of the cursor one character. At the right margin, an automatic new line is performed.

The **addch** subroutine adds the character to the stdscr at the current logical cursor location. To add a character to a user-defined window, use the **waddch** and **mvwaddch** subroutines. The **mvaddch** and **mvwaddch** subroutines move the logical cursor before adding a character.

If you add a character to the bottom of a scrolling region, curses automatically scrolls the region up one line from the bottom of the scrolling region if **scrollok** is enabled. If the character to add is a tab, new-line, or backspace character, curses moves the cursor appropriately in the window to reflect the addition. Tabs are set at every eighth column. If the character is a new-line, curses first uses the **wclrtoeol** subroutine to erase the current line from the logical cursor position to the end of the line before moving the cursor.

You can also use the **addch** subroutines to add control characters to a window. Control characters are drawn in the ^X notation.

#### Adding Video Attributes and Text

Because the *Char* parameter is an integer, not a character, you can combine video attributes with a character by ORing them into the parameter. The video attributes are also set. With this capability you can copy text and video attributes from one location to another using the **inch** (“inch, mvinch, mvwinch, or winch Subroutine” on page 574) and **addch** subroutines.

#### Parameters

*ch*  
*y*  
*x*  
*\*win*

## Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

## Examples

1. To add the character *H* represented by variable *x* to `stdscr` at the current cursor location, enter:

```
chtype x;
x='H';
addch(x);
```

2. To add the *x* character to `stdscr` at the coordinates *y* = 10, *x* = 5, enter:

```
mvaddch(10, 5, 'x');
```

3. To add the *x* character to the user-defined window `my_window` at the coordinates *y* = 10, *x* = 5, enter:

```
WINDOW *my_window;
mvwaddch(my_window, 10, 5, 'x');
```

4. To add the *x* character to the user-defined window `my_window` at the current cursor location, enter:

```
WINDOW *my_window;
waddch(my_window, 'x');
```

5. To add the character *x* in standout mode, enter:

```
waddch(my_window, 'x' | A_STANDOUT);
```

This allows 'x' to be highlighted, but leaves the rest of the window alone.

## Related Information

The **inch**, **winch**, **mvinch**, or **mvwinch** (“inch, mvinch, mvwinch, or winch Subroutine” on page 574) subroutines, **wclrtoeol** (“clrtoeol or wclrtoeol Subroutine” on page 541) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## addnstr, addstr, mvaddnstr, mvaddstr, mvwaddnstr, mvwaddstr, waddnstr, or waddstr Subroutine

### Purpose

Adds a string of multi-byte characters without rendition to a window and advances the cursor.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>

int addnstr(const char *str,
int n);

int addstr(const char *str);

int mvaddnstr(int y,
int x,
const char *str,
int n);

int mvaddstr(int y,
int x,
const char *str);
```

```

int mvwaddnstr(WINDOW *win,
int y,
int x,
const char *str,
int n);
int mvwaddstr(WINDOW *win,
int y,
int x,
const char *str);
int waddnstr(WINDOW *win,
const char *str,
int n);
int waddstr(WINDOW *win,
const char *str);

```

## Description

These subroutines write the characters of the string *str* on the current or specified window starting at the current or specified position using the background rendition.

These subroutines advance the cursor position, perform special character processing, and perform wrapping.

The **addstr**, **mvaddstr**, **mvwaddstr** and **waddstr** subroutines are similar to calling **mbstowcs** on *str*, and then calling **addwstr**, **mvaddwstr**, **mvwaddwstr**, and **waddwstr**, respectively.

The **addnstr**, **mvaddnstr**, **mvwaddnstr** and **waddnstr** subroutines use at most, *n* bytes from *str*. These subroutines add the entire string when *n* is -1.

## Parameters

<i>Column</i>	Specifies the horizontal position to move the cursor to before adding the string.
<i>Line</i>	Specifies the vertical position to move the cursor to before adding the string.
<i>String</i>	Specifies the string to add.
<i>Window</i>	Specifies the window to add the string to.

## Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

## Examples

1. To add the string represented by *xyz* to the *stdscr* at the current cursor location, enter:

```

char *xyz;
xyz="Hello!";
addstr(xyz);

```

2. To add the "Hit a Key" string to the *stdscr* at the coordinates *y=10, x=5*, enter:

```

mvaddstr(10, 5, "Hit a Key");

```

3. To add the *xyz* string to the user-defined window *my\_window* at the coordinates *y=10, x=5*, enter:

```

mvwaddstr(my_window, 10, 5, "xyz");

```

4. To add the *xyz* string to the user-defined string at the current cursor location, enter:

```

waddstr(my_window, "xyz");

```

## Related Information

The **addch** ("addch, mvaddch, mvwaddch, or waddch Subroutine" on page 523) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## attroff, attron, attrset, wattroff, wattron, or wattrset Subroutine

### Purpose

Restricted window attribute control functions.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int attroff (int *attrs);
```

```
int attron (int *attrs);
```

```
int attrset (int *attrs);
```

```
int wattroff (WINDOW *win, int *attrs);
```

```
int wattron (WINDOW *win, int *attrs);
```

```
int wattrset (WINDOW *win, int *attrs);
```

### Description

These subroutines manipulate the window attributes of the current or specified window.

The **attroff** and **wattroff** subroutines turn off *attrs* in the current or specified specified window without affecting any others.

The **attron** and **wattron** subroutines turn on *attrs* in the current or specified specified window without affecting any others.

The **attrset** and **wattrset** subroutines set the background attributes of the current or specified window to *attrs*.

It unspecified whether these subroutines can be used to manipulate attributes than A\_BLINK, A\_BOLD, A\_DIM, A\_REVERSE, A\_STANDOUT and A\_UNDERLINE.

## Parameters

- \*attrs* Specifies which attributes to turn off.
- \*win* Specifies the window in which to turn off the specified attributes.

## Return Values

These subroutines always return either OK or 1.

## Examples

For the **attroff** or **wattroff** subroutines:

1. To turn the off underlining attribute in `stdscr`, enter:  

```
attroff(A_UNDERLINE);
```
2. To turn off the underlining attribute in the user-defined window `my_window`, enter:  

```
wattroff(my_window, A_UNDERLINE);
```

For the **attron** or **wattron** subroutines:

1. To turn on the underlining attribute in `stdscr`, enter:  

```
attron(A_UNDERLINE);
```
2. To turn on the underlining attribute in the user-defined window `my_window`, enter:  

```
wattron(my_window, A_UNDERLINE);
```

For the **attrset** or **wattrset** subroutines:

1. To set the current attribute in the **stdscr** global variable to blink, enter:  

```
attrset(A_BLINK);
```
2. To set the current attribute in the user-defined window `my_window` to blinking, enter:  

```
wattrset(my_window, A_BLINK);
```
3. To turn off all attributes in the **stdscr** global variable, enter:  

```
attrset(0);
```
4. To turn off all attributes in the user-defined window `my_window`, enter:  

```
wattrset(my_window, 0);
```

## Related Information

The **standend** ("standend, standout, wstandend, or wstandout Subroutine" on page 635) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Setting Video Attributes and Curses Options in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## attron or wattron Subroutine

### Purpose

Turns on specified attributes.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
attron( Attributes)  
char *Attributes;
```

```
wattron( Window, Attributes)  
WINDOW *Window;  
char *Attributes;
```

### Description

The **attron** and **wattron** subroutines turn on specified attributes without affecting any others. The **attron** subroutine turns the specified attributes on in stdscr. The **wattron** subroutine turns the specified attributes on in the specified window.

### Parameters

<i>Attributes</i>	Specifies which attributes to turn on.
<i>Window</i>	Specifies the window in which to turn on the specified attributes.

### Examples

1. To turn on the underlining attribute in stdscr, enter:  

```
attron(A_UNDERLINE);
```
2. To turn on the underlining attribute in the user-defined window `my_window`, enter:  

```
wattron(my_window, A_UNDERLINE);
```

### Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Setting Video Attributes and Curses Options in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## attrset or wattrset Subroutine

### Purpose

Sets the current attributes of a window to the specified attributes.

## Libraries

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>
```

```
attrset( Attributes)  
char *Attributes;  
wattrset( Window, Attributes)  
WINDOW *Window;  
char *Attributes;
```

## Description

The **attrset** and **wattrset** subroutines set the current attributes of a window to the specified attributes. The **attrset** subroutine sets the current attribute of **stdscr**. The **wattrset** subroutine sets the current attribute of the specified window.

## Parameters

<i>Attributes</i>	Specifies which attributes to set.
<i>Window</i>	Specifies the window in which to set the attributes.

## Examples

1. To set the current attribute in the **stdscr** global variable to blink, enter:  

```
attrset(A_BLINK);
```
2. To set the current attribute in the user-defined window `my_window` to blinking, enter:  

```
wattrset(my_window, A_BLINK);
```
3. To turn off all attributes in the **stdscr** global variable, enter:  

```
attrset(0);
```
4. To turn off all attributes in the user-defined window `my_window`, enter:  

```
wattrset(my_window, 0);
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Setting Video Attributes and Curses Options in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## baudrate Subroutine

### Purpose

Gets the terminal baud rate.

### Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>

int baudrate(void)
```

## Description

The **baudrate** subroutine extracts the output speed of the terminal in bits per second.

## Return Values

The **baudrate** subroutine returns the output speed of the terminal.

## Examples

To query the baud rate and place the value in the user-defined integer variable BaudRate, enter:  
BaudRate = baudrate();

## Related Information

The **tcgetattr** (“tcgetattr Subroutine” on page 346) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## beep Subroutine

### Purpose

Sounds the audible alarm on the terminal.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>

int beep(void);
```

### Description

The **beep** subroutine alerts the user. It sounds the audible alarm on the terminal, or if that is not possible, it flashes the screen (visible bell). If neither signal is possible, nothing happens.

### Return Values

The **beep** subroutine always returns OK.

### Examples

To sound an audible alarm, enter:  
beep();

### Related Information

The **flash** (“flash Subroutine” on page 557) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Setting Video Attributes and Curses Options in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## box Subroutine

### Purpose

Draws borders from single-byte characters and renditions.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int box(WINDOW *win,  
        chtype verch,  
        chtype horch);
```

### Description

The **box** subroutine draws a border around the edges of the specified window. This subroutine does not advance the cursor position. This subroutine does not perform special character processing or perform wrapping.

The **box** subroutine (*\*win*, *verch*, *horch*) has an effect equivalent to:

```
wborder(win, verch, verch, horch, horch, 0, 0, 0, 0);
```

### Parameters

<i>horch</i>	Specifies the character to draw the horizontal lines of the box. The character must be a 1-column character.
<i>verch</i>	Specifies the character to draw the vertical lines of the box. The character must be a 1-column character.
<i>*win</i>	Specifies the window to draw the box in or around.

### Return Values

Upon successful completion, the **box** function returns OK. Otherwise, it returns ERR.

### Examples

1. To draw a box around the user-defined window, *my\_window*, using | (pipe) as the vertical character and - (minus sign) as the horizontal character, enter:

```
WINDOW *my_window;  
box(my_window, '|', '-');
```

2. To draw a box around *my\_window* using the default characters ACS\_VLINE and ACS\_HLINE, enter:

```
WINDOW *my_window;  
box(my_window, 0, 0);
```

## Related Information

Curses Overview for Programming, List of Curses Subroutines, and Windows in the Curses Environment in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## can\_change\_color, color\_content, has\_colors, init\_color, init\_pair, start\_color or pair\_content Subroutine

### Purpose

Color manipulation functions and external variables for color support.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
bool can_change_color(void);
```

```
int color_content(short color,  
short *red,  
short *green,  
short *blue);
```

```
int COLOR_PAIR(int n);
```

```
bool has_colors(void);
```

```
int init_color  
(short color,  
short red,  
short green,  
short blue);
```

```
int init_pair  
(short pair,  
short f,  
short b);
```

```
int pair_content  
(short pair,  
short *f,  
short *b);
```

```
int PAIR_NUMBER  
(int value);  
int start_color  
(void);
```

```
extern int COLOR_PAIRS;  
extern int COLORS;
```

### Description

These functions manipulate color on terminals that support color.

### Querying Capabilities

The **has\_colors** subroutine indicates whether the terminal is a color terminal. The **can\_change\_color** subroutine indicates whether the terminal is a color terminal on which colors can be redefined.

## Initialisation

The **start\_color** subroutine must be called in order to enable use of colors and before any color manipulation function is called. This subroutine initializes eight basic colors (black, blue, green, cyan, red, magenta, yellow, and white) that can be specified by the color macros (such as **COLOR\_BLACK**) defined in `<curses.h>`. The initial appearance of these eight colors is not specified.

The function also initialises two global external variables:

- **COLORS** defines the number of colors that the terminal supports. If **COLORS** is 0, the terminal does not support redefinition of colors (and **can\_change\_color** subroutine will return **FALSE**).
- **COLOR\_PAIRS** defines the maximum number of color-pairs that the terminal supports.

## Color Identification

The **init\_color** subroutine redefines color number *color*, on terminals that support the redefinition of colors, to have the red, green, and blue intensity components specified by *red*, *green*, and *blue*, respectively. Calling **init\_color** subroutine also changes all occurrences of the specified color on the screen to the new definition.

The **color\_content** subroutine identifies the intensity components of color number *color*. It stores the red, green, and blue intensity components of this color in the addresses pointed to by *red*, *green*, and *blue*, respectively.

For both functions, the color argument must be in the range from **0** to and including **COLORS - 1**. Valid intensity values range from **0** (no intensity component) up to and including **1000** (maximum intensity in that component).

## User-Defined Color Pairs

Calling **init\_pair** defines or redefines color-pair number *pair* to have foreground color *f* and background color *b*. Calling **init\_pair** changes any characters that were displayed in the color pair's old definition to the new definition and refreshes the screen.

After defining the color pair, the macro **COLOR\_PAIR(n)** returns the value of color pair *n*. This value is the color attribute as it would be extracted from a **chtype**. Conversely, the macro **PAIR\_NUMBER(value)** returns the color pair number associated with the color attribute value.

The **pair\_content** subroutine retrieves the component colors of a color-pair number *pair*. It stores the foreground and background color numbers in the variables pointed to by *f* and *b*, respectively.

With **init\_pair** and **pair\_content** subroutines, the value of *pair* must be in a range from **0** to and including **COLOR\_PAIRS - 1**. (There may be an implementation-specific upper limit on the valid value of *pair*, but any such limit is at least 63.) Valid values for *f* and *b* are the range from **0** to and including **COLORS - 1**.

The **can\_change\_color** subroutine returns **TRUE** if the terminal supports colors and can change their definitions; otherwise, it returns **FALSE**.

## Parameters

*color*

*\*red*

*\*green*

*\*blue*

*pair*

*f*

*b*

*value*

## Return Values

The **has\_colors** subroutine returns TRUE if the terminal can manipulate colors; otherwise, it returns FALSE.

Upon successful completion, the other functions return OK. Otherwise, they return ERR.

## Examples

For the **can\_change\_color** subroutine:

To test whether or not a terminal can change its colors, enter the following and check the return for TRUE or FALSE:

```
can_change_color();
```

For the **color\_content** subroutine:

To obtain the RGB component information for color 10 (assuming the terminal supports at least 11 colors), use:

```
short *r, *g, *b;  
color_content(10,r,g,b);
```

For the **has\_color** subroutine:

To determine whether or not a terminal supports color, use:

```
has_colors();
```

For the **pair\_content** subroutine:

To obtain the foreground and background colors for color-pair 5, use:

```
short *f, *b;  
pair_content(5,f,b);
```

For this subroutine to succeed, you must have already initialized the color pair. The foreground and background colors will be stored at the locations pointed to by *f* and *b*.

For the **start\_color** subroutine:

To enable the color support for a terminal that supports color, use:

```
start_color();
```

For the **init\_pair** subroutine:

To initialize the color definition for color-pair 2 to a black foreground (color 0) with a cyan background (color 3), use:

```
init_pair(2,COLOR_BLACK, COLOR_CYAN);
```

For the **init\_color** subroutine:

To initialize the color definition for color 11 to violet on a terminal that supports at least 12 colors, use:

```
init_color(11,500,0,500);
```

## Related Information

The **attroff** (“attroff, attron, attrset, wattroff, wattron, or wattrset Subroutine” on page 526) subroutine.

## **cbreak, nocbreak, noraw, or raw Subroutine**

### **Purpose**

Puts the terminal into or out of CBREAK mode.

### **Library**

Curses Library (**libcurses.a**)

### **Syntax**

```
#include <curses.h>
int cbreak(void);
int nocbreak(void);
int noraw(void);
int raw(void);
```

### **Description**

The **cbreak** subroutine sets the input mode for the current terminal to cbreak mode and overrides a call to the **raw** subroutine.

The **nocbreak** subroutine sets the input mode for the current terminal to Cooked Mode without changing the state of the **ISIG** and **IXON** flags.

The **noraw** subroutine sets the input mode for the current terminal to Cooked Mode and sets the **ISIG** and **IXON** flags.

The **raw** subroutine sets the input mode for the current terminal to Raw Mode.

### **Return Values**

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

### **Examples**

For the **cbreak** and **nocbreak** subroutines:

1. To put the terminal into CBREAK mode, enter:  
`cbreak();`
2. To take the terminal out of CBREAK mode, enter:  
`nocbreak();`
3. To place the terminal into raw mode, use:  
`raw();`
4. To place the terminal out of raw mode, use:  
`noraw();`

For the **noraw** and **raw** subroutines:

1. To place the terminal into raw mode, use:  
`raw();`
2. To place the terminal out of raw mode, use:  
`noraw();`

## Related Information

The **getch** (“getch, mvgetch, mvwgetch, or wgetch Subroutine” on page 561) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## clear, erase, wclear or werase Subroutine

### Purpose

Clears a window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
int clear(void);
int erase(void);

int wclear(WINDOW *win);
int werase(WINDOW *win);
```

### Description

The **clear**, **erase**, **wclear**, and **werase** subroutines clear every position in the current or specified window.

The **clear** and **wclear** subroutines also achieve the same effect as calling the **clearok** subroutine, so that the window is cleared completely on the next call to the **wrefresh** subroutine for the window and is redrawn in its entirety.

### Parameters

*\*win* Specifies the window to clear.

### Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

### Examples

For the **clear** and **wclear** subroutines:

1. To clear stdscr and set a clear flag for the next call to the **refresh** subroutine, enter:  
clear();
2. To clear the user-defined window `my_window` and set a clear flag for the next call to the **wrefresh** subroutine, enter:

```
WINDOW *my_window;
wclear(my_window);
waddstr (my_window, "This will be cleared.");
wrefresh (my_window);
```

3. To erase the standard screen structure, enter:

```
erase();
```

4. To erase the user-defined window `my_window`, enter:

```
WINDOW *my_window;
werase (my_window);
```

**Note:** After the **wrefresh**, the window will be cleared completely. You will not see the string "This will be cleared."

For the **erase** and **werase** subroutines:

1. To erase the standard screen structure, enter:

```
erase();
```

2. To erase the user-defined window `my_window`, enter:

```
WINDOW *my_window;
werase(my_window);
```

## Related Information

The **doupdate** ("doupdate, refresh, wnoutrefresh, or wrefresh Subroutines" on page 657) subroutine, **erase** ("erase or werase Subroutine" on page 555) and **werase** ("erase or werase Subroutine" on page 555) subroutines, **clearok** ("clearok, idlok, leaveok, scrollok, setscreg or wsetscreg Subroutine") subroutine, **refresh** ("refresh or wrefresh Subroutine" on page 608) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## clearok, idlok, leaveok, scrollok, setscreg or wsetscreg Subroutine

### Purpose

Terminal output control subroutines.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int clearok(WINDOW *win,
bool bf);
```

```
int idlok(WINDOW *win,
bool bf);
```

```
int leaveok(WINDOW *win,
bool bf);
```

```
int scrollok(WINDOW *win,
bool bf);
```

```
int setscreg(int top,
int bot);
```

```
int wsetscrreg(WINDOW *win,  
int top,  
int bot);
```

## Description

These subroutines set options that deal with output within Curses.

The **clearok** subroutine assigns the value of *bf* to an internal flag in the specified window that governs clearing of the screen during a refresh. If, during a refresh operation on the specified window, the flag in **curscr** is TRUE or the flag in the specified window is TRUE, then the implementation clears the screen, redraws it in its entirety, and sets the flag to FALSE in **curscr** and in the specified window. The initial state is unspecified.

The **idlok** subroutine specifies whether the implementation may use the hardware insert-line, delete-line, and scroll features of terminals so equipped. If *bf* is TRUE, use of these features is enabled. If *bf* is FALSE, use of these features is disabled and lines are instead redrawn as required. The initial state is FALSE.

The **leaveok** subroutine controls the cursor position after a refresh operation. If *bf* is TRUE, refresh operations on the specified window may leave the terminal's cursor at an arbitrary position. If *bf* is FALSE, then at the end of any refresh operation, the terminal's cursor is positioned at the cursor position contained in the specified window. The initial state is FALSE.

The **scrollok** subroutine controls the use of scrolling. If *bf* is TRUE, then scrolling is enabled for the specified window, with the consequences discussed in Truncation, Wrapping and Scrolling on page 28. If *bf* is FALSE, scrolling is disabled for the specified window. The initial state is FALSE.

The **setscrreg** and **wsetscrreg** subroutines define a software scrolling region in the current or specified window. The *top* and *bot* arguments are the line numbers of the first and last line defining the scrolling region. (Line 0 is the top line of the window.) If this option and the **scrollok** subroutine are enabled, an attempt to move off the last line of the margin causes all lines in the scrolling region to scroll one line in the direction of the first line. Only characters in the window are scrolled. If a software scrolling region is set and the **scrollok** subroutine is not enabled, an attempt to move off the last line of the margin does not reposition any lines in the scrolling region.

## Parameters

The parameters for the **clearok** subroutine are:

<i>Flag</i>	Sets the window clear flag. If TRUE, curses clears the window on the next call to the <b>wrefresh</b> or <b>refresh</b> subroutines. If FALSE, curses does not clear the window.
<i>Window</i>	Specifies the window to clear.

The parameters for the **idlok** subroutine are:

<i>Flag</i>	Specifies whether to enable curses to use the hardware insert/delete line feature (TRUE) or not (FALSE).
<i>Window</i>	Specifies the window it will affect.

The parameters for the **leaveok** subroutine are:

<i>Flag</i>	Specifies whether to leave the physical cursor alone after a refresh (TRUE) or to move the physical cursor to the logical cursor after a refresh (FALSE).
<i>Window</i>	Specifies the window for which to set the <i>Flag</i> parameter.

The parameters for the **scrollok** subroutine are:

*Flag* Enables scrolling when set to TRUE. Otherwise, set the *Flag* parameter to FALSE to disable scrolling.  
*Window* Identifies the window in which to enable or disable scrolling.

The parameters for the **setscrreg** and **wsetscrreg** subroutines are:

*Bmargin* Specifies the last line number in the scrolling region.  
*Tmargin* Specifies the first line number in the scrolling region (0 is the top line of the window).  
*Window* Specifies the window in which to place the scrolling region. You specify this parameter only with the **wsetscrreg** subroutine.

## Return Values

Upon successful completion, the **setscrreg** and **wsetscrreg** subroutines return OK. Otherwise, they return ERR.

The other subroutines always return OK.

## Examples

Examples for the **clearok** subroutine are:

1. To set the user-defined screen `my_screen` to clear on the next call to the **wrefresh** subroutine, enter:

```
WINDOW *my_screen;  
clearok(my_screen, TRUE);
```

2. To set the standard screen structure to clear on the next call to the **refresh** subroutine, enter:

```
clearok(stdscr, TRUE);
```

Examples for the **idlok** subroutine are:

1. To enable curses to use the hardware insert/delete line feature in `stdscr`, enter:

```
idlok(stdscr, TRUE);
```

2. To force curses not to use the hardware insert/delete line feature in the user-defined window `my_window`, enter:

```
idlok(my_window, FALSE);
```

Examples for the **leaveok** subroutine are:

1. To move the physical cursor to the same location as the logical cursor after refreshing the user-defined window `my_window`, enter:

```
WINDOW *my_window;  
leaveok(my_window, FALSE);
```

2. To leave the physical cursor alone after refreshing the user-defined window `my_window`, enter:

```
WINDOW *my_window;  
leaveok(my_window, TRUE);
```

Examples for the **scrollok** subroutine are:

1. To turn scrolling on in the user-defined window `my_window`, enter:

```
WINDOW *my_window;  
scrollok(my_window, TRUE);
```

2. To turn scrolling off in the user-defined window `my_window`, enter:

```
WINDOW *my_window;  
scrollok(my_window, FALSE);
```

Examples for the **setscrreg** or **wsetscrreg** subroutine are:

1. To set a scrolling region starting at the 10th line and ending at the 30th line in the stdscr, enter:  
`setscreg(9, 29);`

**Note:** Zero is always the first line.

2. To set a scrolling region starting at the 10th line and ending at the 30th line in the user-defined window `my_window`, enter:

```
WINDOW *my_window;
wsetscreg(my_window, 9, 29);
```

## Related Information

The **douupdate** (“douupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine, **scr** (“scr, scroll, wscr Subroutine” on page 619) subroutine, **refresh** or **wrefresh** (“refresh or wrefresh Subroutine” on page 608) subroutine.

Curses Library, List of Additional Curses Subroutines, and Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## clrrobot or wclrrobot Subroutine

### Purpose

Erases the current line from the logical cursor position to the end of the window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int clrrobot(void);
```

```
int wclrrobot(WINDOW *win);
```

### Description

The **clrrobot** and **wclrrobot** subroutines erase all lines following the cursor in the current or specified window, and erase the current line from the cursor to the end of the line, inclusive. These subroutines do not update the cursor.

### Parameters

*\*win* Specifies the window in which to erase lines.

### Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

### Examples

1. To erase the lines below and to the right of the logical cursor in the stdscr, enter:

```
clrrobot();
```

2. To erase the lines below and to the right of the logical cursor in the user-defined window `my_window`, enter:

```
WINDOW *my_window;
wclrrobot(my_window);
```

## Related Information

The **doudate** (“doudate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## clrtoeol or wclrtoeol Subroutine

### Purpose

Erases the current line from the logical cursor position to the end of the line.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
int clrtoeol(void);

int wclrtoeol(WINDOW * win);
```

### Description

The **clrtoeol** and **wclrtoeol** subroutines erase the current line from the cursor to the end of the line, inclusive, in the current or specified window. These subroutines do not update the cursor.

### Parameters

*\*win* Specifies the window in which to clear the line.

### Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

### Examples

1. To clear the line to the right of the logical cursor in the stdscr, enter:  

```
clrtoeol();
```
2. To clear the line to the right of the logical cursor in the user-defined window `my_window`, enter:  

```
WINDOW *my_window;
wclrtoeol(my_window);
```

## Related Information

The **doudate** (“doudate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## color\_content Subroutine

### Purpose

Returns the current intensity of the red, green, and blue (RGB) components of a color.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
color_content(Color, R, G,
B)
short Color;
short *R, *G, *B;
```

### Description

The **color\_content** subroutine, given a color number, returns the current intensity of its red, green, and blue (RGB) components. This subroutine stores the information in the address specified by the *R*, *G*, and *B* arguments. If successful, this returns OK. Otherwise, this subroutine returns ERR if the color does not exist, is outside the valid range, or the terminal cannot change its color definitions.

To determine if you can change the color definitions for a terminal, use the **can\_change\_color** subroutine. You must call the **start\_color** subroutine before you can call the **color\_content** subroutine.

**Note:** The values stored at the addresses pointed to by *R*, *G*, and *B* are between 0 (no component) and 1000 (maximum amount of component) inclusive.

### Return Values

**OK** Indicates the subroutine was successful.

**ERR** Indicates the color does not exist, is outside the valid range, or the terminal cannot change its color definitions.

### Parameters

*B* Points to the address that stores the intensity value of the blue component.  
*Color* Specifies the color number. The color parameter must be a value between **0** and **COLORS-1** inclusive.  
*R* Points to the address that stores the intensity value of the red component.  
*G* Points to the address that stores the intensity value of the green component.

### Example

To obtain the RGB component information for color 10 (assuming the terminal supports at least 11 colors), use:

```
short *r, *g, *b; color_content(10,r,g,b);
```

## Related Information

The **start\_color** (“start\_color Subroutine” on page 636) subroutine.

Curses Overview for Programming, Manipulating Video Attributes,

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## copywin Subroutine

### Purpose

Copies a region of a window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int copywin(const WINDOW *srcwin,  
WINDOW *dstwin,  
int sminrow,  
int smincol,  
int dminrow,  
int dmincol,  
int dmaxrow,  
int dmaxcol,  
int overlay);
```

### Description

The **copywin** subroutine provides a finer granularity of control over the **overlay** and **overwrite** subroutines. As in the **prefresh** subroutine, a rectangle is specified in the destination window, (*dminrow*, *dmincol*) and (*dmaxrow*, *dmaxcol*), and the upper-left-corner coordinates of the source window, (*sminrow*, *smincol*). If the **overlay** subroutine is TRUE, then copying is non-destructive, as in the **overlay** subroutine. If the **overlay** subroutine is FALSE, then copying is destructive, as in the **overwrite** subroutine.

### Parameters

<i>*srcwin</i>	Points to the source window containing the region to copy.
<i>*dstwin</i>	Points to the destination window to copy into.
<i>sminrow</i>	Specifies the upper left row coordinate of the source region.
<i>smincol</i>	Specifies the upper left column coordinate of the source region.
<i>dminrow</i>	Specifies the upper left row coordinate of the destination region.
<i>dmincol</i>	Specifies the upper left column coordinate for the destination region.
<i>dmaxrow</i>	Specifies the lower right row coordinate for the destination region.
<i>dmaxcol</i>	Specifies the lower right column coordinate for the destination region.
<i>overlay</i>	Sets the type of copy. If set to TRUE the copy is nondestructive. Otherwise, if set to FALSE, the copy is destructive.

### Return Values

Upon successful completion, the **copywin** subroutine returns OK. Otherwise, it returns ERR.

## Examples

To copy to an area in the destination window defined by coordinates (30,40), (30,49), (39,40), and (39,49) beginning with coordinates (0,0) in the source window, enter the following:

```
WINDOW *srcwin, *dstwin;
```

```
copywin(srcwin, dstwin,  
0, 0, 30,40, 39, 49,  
TRUE);
```

The example copies ten rows and ten columns from the source window beginning with coordinates (0,0) to the region in the destination window defined by the upper left coordinates (30, 40) and lower right coordinates (39, 49). Because the `Overlay` parameter is set to `TRUE`, the copy is nondestructive and blanks from the source window are not copied.

## Related Information

The **newpad** (“newpad, pnoutrefresh, prefresh, or subpad Subroutine” on page 592) and **overlay or overwrite** (“overlay or overwrite Subroutine” on page 601) subroutines.

Curses Overview for Programming, Manipulating Window Data with Curses Manipulating Characters with Curses, List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*

---

## curs\_set Subroutine

### Purpose

Sets the cursor visibility.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int curs_set(int visibility);
```

### Description

The **curs\_set** subroutine sets the appearance of the cursor based on the value of visibility:

Value of visibility	Appearance of Cursor
---------------------	----------------------

- |   |  |
|---|--|
| 0 | invisible                              |
| 1 | terminal-specific normal mode          |
| 2 | terminal-specific high visibility mode |

The terminal does not necessarily support all the above values.

## Parameters

*Visibility* Sets the cursor state. You can set the cursor state to one of the following:

<b>0</b>	Invisible
<b>1</b>	Visible
<b>2</b>	Very visible

## Return Values

If the terminal supports the cursor mode specified by *visibility*, then the **cur\_set** subroutine returns the previous cursor state. Otherwise, the subroutine returns ERR.

## Examples

To set the cursor state to invisible, use:

```
cur_set(0);
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*

Setting Video Attributes in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*

---

## def\_prog\_mode, def\_shell\_mode, reset\_prog\_mode or reset\_shell\_mode Subroutine

### Purpose

Saves/restores the program or shell terminal modes.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int def_prog_mode  
(void);
```

```
int def_shell_mode  
(void);
```

```
int reset_prog_mode  
(void);
```

```
int reset_shell_mode  
(void);
```

## Description

The **def\_prog\_mode** subroutine saves the current terminal modes as the "program" (in Curses) state for use by the **reset\_prog\_mode** subroutine.

The **def\_shell\_mode** subroutine saves the current terminal modes as the "shell" (not in Curses) state for use by the **reset\_shell\_mode** subroutine.

The **reset\_prog\_mode** subroutine restores the terminal to the "program" (in Curses) state.

The **reset\_shell\_mode** subroutine restores the terminal to the "shell" (not in Curses) state.

These subroutines affect the mode of the terminal associated with the current screen.

## Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

## Examples

For the **def\_prog\_mode** subroutine:

To save the "in curses" state, enter:

```
def_prog_mode();
```

For the **def\_shell\_mode** subroutine:

To save the "out of curses" state, enter:

```
def_shell_mode();
```

This routine saves the "out of curses" state.

## Related Information

The **doupdate** ("doupdate, refresh, wnoutrefresh, or wrefresh Subroutines" on page 657), **endwin** ("endwin Subroutine" on page 554), **initscr** ("initscr and newterm Subroutine" on page 577), and the **setupterm** ("setupterm Subroutine" on page 624) subroutines.

Curses Overview for Programming, List of Curses Subroutines, Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## def\_shell\_mode Subroutine

### Purpose

Saves the current terminal modes as shell mode ("out of curses").

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
def_shell_mode( )
```

## Description

The **def\_shell\_mode** subroutine saves the current terminal driver line discipline modes in the current terminal structure for later use by **reset\_shell\_mode()**. The **def\_shell\_mode** subroutine is called automatically by the **setupterm** subroutine.

This routine would normally not be called except by a library routine.

## Example

To save the "out of curses" state, enter:

```
def_shell_mode();
```

This routine saves the "out of curses" state.

## Related Information

The **setupterm** ("setupterm Subroutine" on page 624) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## del\_curterm, restartterm, set\_curterm, or setupterm Subroutine

### Purpose

Interfaces to the **terminfo** database.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <term.h>
```

```
int del_curterm(TERMINAL *oterm);
```

```
int restartterm(char *term,  
int fildes,  
int *erret);
```

```
TERMINAL *set_curterm(TERMINAL *nterm);
```

```
int setupterm(char *term,  
int fildes,  
int *erret);
```

### Description

The **del\_curterm**, **restartterm**, **set\_curterm**, **setupterm** subroutines retrieve information from the **terminfo** database.

To gain access to the **terminfo** database, the **setupterm** subroutine must be called first. It is automatically called by the **initscr** and **newterm** subroutines. The **setupterm** subroutine initialises the other subroutines to use the **terminfo** record for a specified terminal (which depends on whether the **use\_env** subroutine was called). It sets the **dur\_term** external variable to a **TERMINAL** structure that contains the record from the **terminfo** database for the specified terminal.

The terminal type is the character string **term**; if **term** is a null pointer, the environment variable **TERM** is used. If **TERM** is not set or if its value is an empty string, the "unknown" is used as the terminal type. The

application must set the *fildev* parameter to a file descriptor, open for output, to the terminal device, before calling the **setupterm** subroutine. If the *erret* parameter is not null, the integer it points to is set to one of the following values to report the function outcome:

- 1 The terminfo database was not found (function fails).
- 0 The entry for the terminal was not found in **terminfo** (function fails).
- 1 Success.

A simple call to the **setupterm** subroutine that uses all the defaults and sends the output to stdout is:  
`setupterm(char *)0, fileno(stdout), (int *)0);`

The **set\_curterm** subroutine sets the variable **cur\_term** to *nterm*, and makes all of the terminfo boolean, numeric, and string variables use the values from *nterm*.

The **del\_curterm** subroutine frees the space pointed to by *oterm* and makes it available for further use. If *oterm* is the same as **cur\_term**, references to any of the terminfo boolean, numeric, and string variables thereafter may refer to invalid memory locations until the **setupterm** subroutine is called again.

The **restartterm** subroutine assumes a previous call to the **setupterm** subroutine (perhaps from the **initscr** or **newterm** subroutine). It lets the application specify a different terminal type in *term* and updates the information returned by the **baudrate** subroutine based on the *fildev* parameter, but does not destroy other information created by the **initscr**, **newterm**, or **setupterm** subroutines.

## Parameters

*\*oterm*  
*\*term*  
*fildev*  
*\*erret*  
*\*nterm*

## Return Values

Upon successful completion, the **set\_curterm** subroutine returns the previous value of **cur\_term**. Otherwise, it returns a null pointer.

Upon successful completion, the other subroutines return OK. Otherwise, they return ERR.

## Examples

To free the space occupied by a **TERMINAL** structure called *my\_term*, use:

```
TERMINAL *my_term; del_curterm(my_term);
```

For the **restartterm** subroutine:

To restart an  **aixterm** after a previous memory save and exit on error with a message, enter:  
`restartterm("aixterm", 1, (int*)0);`

For the **set\_curterm** subroutine:

To set the **cur\_term** variable to point to the *my\_term* terminal, use:

```
TERMINAL *newterm; set_curterm(newterm);
```

For the **setupterm** subroutine:

To determine the current terminal's capabilities using **\$TERM** as the terminal name, standard output as output, and returning no error codes, enter:

```
setupterm((char*) 0, 1, (int*) 0);
```

## Related Information

The **baudrate** (“baudrate Subroutine” on page 529) subroutine, **longname** (“longname Subroutine” on page 586) subroutine, **putc** subroutine, **tgetent** (“tgetent, tgetflag, tgetnum, tgetstr, or tgoto Subroutine” on page 640) subroutine, **tigetflag** (“tigetflag, tigetnum, tigetstr, or tparm Subroutine” on page 644) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## delay\_output Subroutine

### Purpose

Sets the delay output.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int delay_output(int ms);
```

### Description

On terminals that support pad characters, the **delay\_output** subroutine pauses the output for at least *ms* milliseconds. Otherwise, the length of the delay is unspecified.

### Parameters

*ms* Specifies the number of milliseconds to delay output.

### Return Values

Upon successful completion, the **delay\_output** subroutine returns OK. Otherwise, it returns ERR.

### Examples

To set the output to delay 250 milliseconds, enter:

```
delay_output(250);
```

## Related Information

Curses Overview for Programming, List of Curses Subroutines, Understanding Terminals with Curses in AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs.

---

## delch, mvdelch, mvwdelch or wdelch Subroutine

### Purpose

Deletes the character from a window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int delch(void);
```

```
int mvdelch  
(int y  
int x);
```

```
mvwdelch  
(WINDOW *win;  
int y  
int x);
```

```
wdelch  
(WINDOW *win);
```

### Description

The **delch**, **mvdelch**, **mvwdelch**, and **wdelch** subroutines delete the character at the current or specified position in the current or specified window. This subroutine does not change the cursor position.

### Parameters

*x*  
*y*  
*\*win* Identifies the window from which to delete the character.

### Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

### Examples

1. To delete the character at the current cursor location in the standard screen structure, enter:  
`mvdelch();`
2. To delete the character at cursor position *y*=20 and *x*=30 in the standard screen structure, enter:  
`mvwdelch(20, 30);`
3. To delete the character at cursor position *y*=20 and *x*=30 in the user-defined window *my\_window*, enter:  
`wdelch(my_window, 20, 30);`

## Related Information

Curses Overview for Programming, List of Curses Subroutines, Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## deleteln or wdeleteln Subroutine

### Purpose

Deletes lines in a window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
int deleteln(void);

int wdeleteln(WINDOW *win);
```

### Description

The **deleteln** and **wdeleteln** subroutines delete the line containing the cursor in the current or specified window and move all lines following the current line one line toward the cursor. The last line of the window is cleared. The cursor position does not change.

### Parameters

*\*win* Specifies the window in which to delete the line.

### Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

### Examples

1. To delete the current line in stdscr, enter:  

```
deleteln();
```
2. To delete the current line in the user-defined window `my_window`, enter:  

```
WINDOW *my_window;
wdeleteln(my_window);
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## delwin Subroutine

### Purpose

Deletes a window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int delwin(WINDOW *win);
```

### Description

The **delwin** subroutine deletes *win*, freeing all memory associated with it. The application must delete subwindows before deleting the main window.

### Parameters

*\*win* Specifies the window to delete.

### Return Values

Upon successful completion, the **delwin** subroutine returns OK. Otherwise, it returns ERR.

### Examples

To delete the user-defined window *my\_window* and its subwindow *my\_sub\_window*, enter:

```
WINDOW *my_sub_window, *my_window;  
delwin(my_sub_window);
```

```
delwin(my_window);
```

### Related Information

The **derwin** (“*derwin*, *newwin*, or *subwin* Subroutine” on page 596) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Window Data with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## echo or noecho Subroutine

### Purpose

Enables/disables terminal echo.

### Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>
int echo(void);
int noecho(void);
```

## Description

The **echo** subroutine enables Echo mode for the current screen. The **noecho** subroutine disables Echo mode for the current screen. Initially, curses software echo mode is enabled and hardware echo mode of the tty driver is disabled. The **echo** and **noecho** subroutines control software echo only. Hardware echo must remain disabled for the duration of the application, else the behaviour is undefined.

## Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

## Examples

1. To turn echoing on, use:  
    echo();
2. To turn echoing off, use:  
    noecho();

## Related Information

The **wgetch** (“getch, mvgetch, mvwgetch, or wgetch Subroutine” on page 561) subroutine

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines and Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## echochar or wechochar Subroutines

### Purpose

Echos single-byte character and rendition to a window and refreshes the window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
int echochar(const chtype ch);
int wechochar(WINDOW *win,
const chtype ch);
```

### Description

The **echochar** subroutine is equivalent to a call to the **addch** subroutine followed by a call to the **refresh** subroutine.

The **wechochar** subroutine is equivalent to a call to the **waddch** subroutine followed by a call to the **wrefresh** subroutine.

## Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

## Example

To output the character I to the stdscr at the present cursor location and to update the physical screen, do the following:

```
echochar('I');
```

## Related Information

The **addch**, **doupage**, **echo\_wchar**, **waddch**, **wmvaddch**, and **mvaddch** (“addch, mvaddch, mvwaddch, or waddch Subroutine” on page 523) subroutines.

Curses Overview for Programming and List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## endwin Subroutine

### Purpose

Suspends curses session.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int endwin(void)
```

### Description

The **endwin** subroutine restores the terminal after Curses activity by at least restoring the saved shell terminal mode, flushing any output to the terminal and moving the cursor to the first column of the last line of the screen. Refreshing a window resumes program mode. The application must call the **endwin** subroutine for each terminal being used before exiting. If the **newterm** subroutine is called more than once for the same terminal, the first screen created must be the last one for which the **endwin** subroutine is called.

### Return Values

Upon successful completion, the **endwin** subroutine returns OK. Otherwise, it returns ERR.

### Examples

To terminate curses permanently or temporarily, enter:

```
endwin();
```

## Related Information

The **doupage** (“doupage, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine, **initscr** (“initscr and newterm Subroutine” on page 577) subroutine, and **isendwin** subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Starting and Stopping Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## erase or werase Subroutine

### Purpose

Copies blank spaces to every position in a window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
erase( )
```

```
werase( Window)
WINDOW *Window;
```

### Description

The **erase** and **werase** subroutines copy blank spaces to every position in the specified window. Use the **erase** subroutine with the stdscr and the **werase** subroutine with user-defined windows.

### Parameters

*Window* Specifies the window to erase.

### Examples

1. To erase the standard screen structure, enter:  

```
erase();
```
2. To erase the user-defined window `my_window`, enter:  

```
WINDOW *my_window;
werase(my_window);
```

### Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## erasechar, erasewchar, killchar, and killwchar Subroutine

### Purpose

Terminal environment query functions.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>

char erasechar(void);

int erasewchar(wchar_t *ch);

char killchar(void);

int killwchar(wchar_t
*ch);
```

### Description

The **erasechar** subroutine returns the current character chosen by the user. The **erasechar** subroutine stores the current erase character in the object pointed to by the *ch* parameter. If no erase character has been defined, the subroutine will fail and the object pointed to by *ch* will not be changed.

The **killchar** subroutine returns the current line.

The **killchar** subroutine stores the current line kill character in the object pointed to by *ch*. If no line kill character has been defined, the subroutine will fail and the object pointed to by *ch* will not be changed.

### Return Values

The **erasechar** subroutine returns the erase character and the **killchar** subroutine returns the line kill character. The return value is unspecified when these characters are multi-byte characters.

Upon successful completion, the **erasechar** subroutine and the **killchar** subroutine return OK. Otherwise, they return ERR.

### Examples

To retrieve a user's erase character and return it to the user-defined variable `myerase`, enter:

```
myerase = erasechar();
```

### Related Information

The **clearok** ("clearok, idlok, leaveok, scrollok, setscreg or wsetscreg Subroutine" on page 537) subroutine, **tcgetattr** ("tcgetattr Subroutine" on page 346) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## filter Subroutine

### Purpose

Disables use of certain terminal capabilities.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
void filter(void);
```

### Description

The **filter** subroutine changes the algorithm for initialising terminal capabilities that assume that the terminal has more than one line. A subsequent call to the **initscr** or **newterm** subroutine performs the following actions:

- Disables use of `clear`, `cu`, `cu1`, `cu1`, `cu1`, and `vpa`.
- Sets the value of the home string to the value of the `cr` string.
- Sets lines equal to 1.

Any call to the **filter** subroutine must precede the call to the **initscr** or **newterm** subroutine.

### Related Information

The **initscr** (“`initscr` and `newterm` Subroutine” on page 577) subroutine, **newterm** (“`newterm` Subroutine” on page 594) subroutine.

Curses Overview for Programming and List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## flash Subroutine

### Purpose

Flashes the screen.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
int flash(void);
```

### Description

The **flash** subroutine alerts the user. It flashes the screen, or if that is not possible, it sounds the audible alarm on the terminal. If neither signal is possible, nothing happens.

### Return Values

The **flash** subroutine always returns OK.

## Examples

To cause the terminal to flash, enter:

```
flash();
```

## Related Information

The **beep** (“beep Subroutine” on page 530) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Setting Video Attributes and Curses Options in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## flushinp Subroutine

### Purpose

Discards input.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
int flushinp(void);
```

### Description

The **flushinp** subroutine discards (flushes) any characters in the input buffers associated with the current screen.

### Return Values

The **flushinp** subroutine always returns OK.

## Examples

To flush all type-ahead characters typed by the user but not yet read by the program, enter:

```
flushinp();
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## garbagedlines Subroutine

### Purpose

Discards and replaces a number of lines in a window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
garbagedlines(Window, BegLine, NumLines)
```

```
WINDOW * Window;
```

```
int BegLine, NumLines;
```

### Description

The **garbagedlines** subroutine discards and replaces lines in a window. The *BegLine* parameter specifies the beginning line number and the *NumLines* parameter specifies the number of lines to discard. Curses discards and replaces the specified lines before adding more data.

Uses this subroutine for applications that need to redraw a line that is garbled. Lines may become garbled as the result of noisy communication lines. Instead of refreshing the entire display, use the **garbagedlines** subroutine to refresh a portion of the display and to avoid even more communication noise.

### Parameters

<i>Window</i>	Points to a window.
<i>BegLine</i>	Identifies the beginning line in a range of lines to discard.
<i>NumLines</i>	Specifies the total number of lines in a range of lines to discard and replace.

### Examples

To discard and replace 5 lines in the `mywin` window starting with line 10, use:

```
WINDOW *mywin; garbagedlines(mywin,10, 5);
```

### Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Window Data with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## getbegyx, getmaxyx, getparyx, or getyx Subroutine

### Purpose

Gets the cursor and window coordinates.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
include <curses.h>
```

```
void getbegyx(WINDOW *win,  
int y,  
int x);
```

```
void getmaxyx(WINDOW *win,  
int y,  
int x);
```

```
void getparyx(WINDOW *win,  
int y,  
int x);
```

```
void getyx(WINDOW *win,  
int y,  
int x);
```

### Description

The **getbegyx** macro stores the absolute screen coordinates of the specified window's origin in *y* and *x*.

The **getmaxyx** macro stores the number of rows of the specified window in *y* and *x* and stores the window's number of columns in *x*.

The **getparyx** macro, if the specified window is a subwindow, stores in *y* and *x* the coordinates of the window's origin relative to its parent window. Otherwise, -1 is stored in *y* and *x*.

The **getyx** macro stores the cursor position of the specified window in *y* and *x*.

### Parameters

*\*win*    Identifies the window to get the coordinates from.  
*Y*       Returns the row coordinate.  
*X*       Returns the column coordinate.

### Examples

For the **getbegyx** subroutine:

To obtain the beginning coordinates for the `my_win` window and store in integers *y* and *x*, use:

```
WINDOW *my_win;  
int y, x;  
getbegyx(my_win, y, x);
```

For the **getmaxyx** subroutine:

To obtain the size of the `my_win` window, use:

```
WINDOW *my_win;

int y,x;
getmaxyx(my_win, y, x);
```

Integers `y` and `x` will contain the size of the window.

## Related Information

Controlling the Cursor with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## getch, mvgetch, mvwgetch, or wgetch Subroutine

### Purpose

Gets a single-byte character from the terminal.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>

int getch(void)

int mvgetch(int y,
int x);

int mvwgetch(WINDOW *win,
int y,
int x);

int wgetch(WINDOW *win);
```

### Description

The **getch**, **wgetch**, **mvgetch**, and **mvwgetch** subroutines read a single-byte character from the terminal associated with the current or specified window. The results are unspecified if the input is not a single-byte character. If the  **keypad** subroutine is enabled, these subroutines respond to the corresponding **KEY\_** value defined in `<curses.h>`.

Processing of terminal input is subject to the general rules described in Section 3.5 on page 34.

If echoing is enabled, then the character is echoed as though it were provided as an input argument to the **addch** subroutine, except for the following characters:

```
<backspace>,
<left-arrow> and
the current erase character:
```

The input is interpreted as specified in Section 3.4.3 on page 31 and then the character at the resulting cursor position is deleted as though the **delch** subroutine was called, except that if the cursor was originally in the first column of the line, then the user is alerted as though the **beep** subroutine was called.

The user is alerted as though the **beep** subroutine was called. Information concerning the function keys is not returned to the caller.

## Function Keys

If the current or specified window is not a pad, and it has been moved or modified since the last refresh operation, then it will be refreshed before another character is read.

## The Importance of Terminal Modes

The output of the **getch** subroutines is, in part, determined by the mode of the terminal. The following describes the action of the **getch** subroutines in each type of terminal mode:

Mode	Action of <b>getch</b> Subroutines
<b>NODELAY</b>	Returns a value of ERR if there is no input waiting.
<b>DELAY</b>	Halts execution until the system passes text through the program. If CBREAK mode is also set, the program stops after receiving one character. If <b>NOCBREAK</b> mode is set, the <b>getch</b> subroutine stops reading after the first new line character.
<b>HALF-DELAY</b>	Halts execution until a character is typed or a specified time out is reached. If <b>echo</b> is set, the character is also echoed to the window.

**Note:** When using the **getch** subroutines do not set both the **NOCBREAK** mode and the **ECHO** mode at the same time. This can cause undesirable results depending on the state of the tty driver when each character is typed.

## Getting Function Keys

If your program enables the keyboard with the **keypad** subroutine, and the user presses a function key, the token for that function key is returned instead of raw characters. The possible function keys are defined in the **/usr/include/curses.h** file. Each **#define** macro begins with a **KEY\_** prefix.

If a character is received that could be the beginning of a function key (such as an Escape character) curses sets a timer. If the remainder of the sequence is not received before the timer expires, the character is passed through. Otherwise, the function key's value is returned. For this reason, after a user presses the Esc key there is a delay before the escape is returned to the program. Programmers should not use the Esc key for a single character routine.

Within the **getch** subroutine, a structure of type **timeval**, defined in the **/usr/include/sys/time.h** file, indicates the maximum number of microseconds to wait for the key response to complete.

The **ESCDELAY** environment variable sets the length of time to wait before timing out and treating the ESC keystroke as the ESC character rather than combining it with other characters in the buffer to create a key sequence. The **ESCDELAY** environment variable is measured in fifths of a millisecond. If **ESCDELAY** is 0, the system immediately composes the **ESCAPE** response without waiting for more information from the buffer. The user may choose any value between 0 and 99,999, inclusive. The default setting for the **ESCDELAY** environment variable is 500 (one tenth of a second).

Programs that do not want the **getch** subroutines to set a timer can call the **notimeout** subroutine. If **notimeout** is set to TRUE, curses does not distinguish between function keys and characters when retrieving data.

The **getch** subroutines might not be able to return all function keys because they are not defined in the **terminfo** database or because the terminal does not transmit a unique code when the key is pressed. The following function keys may be returned by the **getch** subroutines:

<b>KEY_MIN</b>	Minimum curses key.
<b>KEY_BREAK</b>	Break key (unreliable).
<b>KEY_DOWN</b>	Down Arrow key.
<b>KEY_UP</b>	Up Arrow key.
<b>KEY_LEFT</b>	Left Arrow key.
<b>KEY_RIGHT</b>	Right Arrow key.
<b>KEY_HOME</b>	Home key.
<b>KEY_BACKSPACE</b>	Backspace.
<b>KEY_F(<i>n</i>)</b>	Function key $F_n$ , where $n$ is an integer from 0 to 64.
<b>KEY_DL</b>	Delete line.
<b>KEY_IL</b>	Insert line.
<b>KEY_DC</b>	Delete character.
<b>KEY_IC</b>	Insert character or enter insert mode.
<b>KEY_EIC</b>	Exit insert character mode.
<b>KEY_CLEAR</b>	Clear screen.
<b>KEY_EOS</b>	Clear to end of screen.
<b>KEY_EOL</b>	Clear to end of line.
<b>KEY_SF</b>	Scroll 1 line forward.
<b>KEY_SR</b>	Scroll 1 line backwards (reverse).
<b>KEY_NPAGE</b>	Next page.
<b>KEY_PPAGE</b>	Previous page.
<b>KEY_STAB</b>	Set tab.
<b>KEY_CTAB</b>	Clear tab.
<b>KEY_CATAB</b>	Clear all tabs.
<b>KEY_ENTER</b>	Enter or send (unreliable).
<b>KEY_SRESET</b>	Soft (partial) reset (unreliable).
<b>KEY_RESET</b>	Reset or hard reset (unreliable).
<b>KEY_PRINT</b>	Print or copy.
<b>KEY_LL</b>	Home down or bottom (lower left).
<b>KEY_A1</b>	Upper-left key of keypad.
<b>KEY_A3</b>	Upper-right key of keypad.
<b>KEY_B2</b>	Center-key of keypad.
<b>KEY_C1</b>	Lower-left key of keypad.
<b>KEY_C3</b>	Lower-right key of keypad.
<b>KEY_BTAB</b>	Back tab key.
<b>KEY_BEG</b>	beg(inning) key
<b>KEY_CANCEL</b>	cancel key
<b>KEY_CLOSE</b>	close key
<b>KEY_COMMAND</b>	cmd (command) key
<b>KEY_COPY</b>	copy key
<b>KEY_CREATE</b>	create key
<b>KEY_END</b>	end key
<b>KEY_EXIT</b>	exit key
<b>KEY_FIND</b>	find key
<b>KEY_HELP</b>	help key
<b>KEY_MARK</b>	mark key
<b>KEY_MESSAGE</b>	message key
<b>KEY_MOVE</b>	move key
<b>KEY_NEXT</b>	next object key
<b>KEY_OPEN</b>	open key
<b>KEY_OPTIONS</b>	options key
<b>KEY_PREVIOUS</b>	previous object key
<b>KEY_REDO</b>	redo key
<b>KEY_REFERENCE</b>	ref(erence) key
<b>KEY_REFRESH</b>	refresh key

<b>KEY_REPLACE</b>	replace key
<b>KEY_RESTART</b>	restart key
<b>KEY_RESUME</b>	resume key
<b>KEY_SAVE</b>	save key
<b>KEY_SBEG</b>	shifted beginning key
<b>KEY_SCANCEL</b>	shifted cancel key
<b>KEY_SCOMMAND</b>	shifted command key
<b>KEY_SCOPY</b>	shifted copy key
<b>KEY_SCREATE</b>	shifted create key
<b>KEY_SDC</b>	shifted delete char key
<b>KEY_SDL</b>	shifted delete line key
<b>KEY_SELECT</b>	select key
<b>KEY_SEND</b>	shifted end key
<b>KEY_SEOL</b>	shifted clear line key
<b>KEY_SEXIT</b>	shifted exit key
<b>KEY_SFIND</b>	shifted find key
<b>KEY_SHELP</b>	shifted help key
<b>KEY_SHOME</b>	shifted home key
<b>KEY_SIC</b>	shifted input key
<b>KEY_SLEFT</b>	shifted left arrow key
<b>KEY_SMESSAGE</b>	shifted message key
<b>KEY_SMOVE</b>	shifted move key
<b>KEY_SNEXT</b>	shifted next key
<b>KEY_SOPTIONS</b>	shifted options key
<b>KEY_SPREVIOUS</b>	shifted prev key
<b>KEY_SPRINT</b>	shifted print key
<b>KEY_SREDO</b>	shifted redo key
<b>KEY_SREPLACE</b>	shifted replace key
<b>KEY_SRIGHT</b>	shifted right arrow
<b>KEY_SRSUME</b>	shifted resume key
<b>KEY_SSAVE</b>	shifted save key
<b>KEY_SSUSPEND</b>	shifted suspend key
<b>KEY_SUNDO</b>	shifted undo key
<b>KEY_SUSPEND</b>	suspend key
<b>KEY_UNDO</b>	undo key

## Parameters

<i>Column</i>	Specifies the horizontal position to move the logical cursor to before getting the character.
<i>Line</i>	Specifies the vertical position to move the logical cursor to before getting the character.
<i>Window</i>	Identifies the window to get the character from and echo it into.

## Return Values

Upon successful completion, the **getch**, **mvwgetch**, and **wgetch** subroutines, CURSES, and Curses Interface return the single-byte character, KEY\_ value, or ERR. When in the nodelay mode and no data is available, ERR is returned.

## Examples

1. To get a character and echo it to the stdscr, use:

```
mvwgetch();
```
2. To get a character and echo it into stdscr at the coordinates y=20, x=30, use:

```
mvwgetch(20, 30);
```

3. To get a character and echo it into the user-defined window `my_window` at coordinates `y=20, x=30`, use:

```
WINDOW *my_window;  
mvwgetch(my_window, 20, 30);
```

## Related Information

The **cbreak** (“cbreak, nocbreak, noraw, or raw Subroutine” on page 535), **douupdate** (“douupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657), and **insch** (“insch, mvinsch, mvwinsch, or winsch Subroutine” on page 578) subroutines, **keypad** (“keypad Subroutine” on page 582) subroutine, **meta** (“meta Subroutine” on page 588) subroutine, **nodelay** (“nodelay Subroutine” on page 599) subroutine, **echo** or **noecho** (“echo or noecho Subroutine” on page 552) subroutine, **notimeout** (“notimeout, timeout, wtimeout Subroutine” on page 599) subroutine, **ebreak** or **nocbreak** (“cbreak, nocbreak, noraw, or raw Subroutine” on page 535) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## getmaxyx Subroutine

### Purpose

Returns the size of a window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
getmaxyx( Window, Y, X);  
WINDOW *Window;  
int Y, X;
```

### Description

The **getmaxyx** subroutine returns the size of a window. The size is returned as the number of rows and columns in the window. The values are stored in integers `Y` and `X`.

### Parameters

<code>Window</code>	Identifies the window whose size to get.
<code>Y</code>	Contains the number of rows in the window.
<code>X</code>	Contains the number of columns in the window.

### Example

To obtain the size of the `my_win` window, use:

```
WINDOW *my_win;

int y,x;
getmaxyx(my_win, y, x);
```

Integers *y* and *x* will contain the size of the window.

## Related Information

Controlling the Cursor with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## getnstr, getstr, mvgetnstr, mvgetstr, mvwgetnstr, mvwgetstr, wgetnstr, or wgetstr Subroutine

### Purpose

Gets a multi-byte character string from the terminal.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int getnstr(char *str,
int n);
```

```
int getstr(char *str);
```

```
int mvgetnstr(int y,
int x,
char *st,
int n);
```

```
int mvgetstr(int y,
int x,
char *str);
```

```
int mvwgetnstr(WINDOW *win,
int y,
int x,
char *str,
int n);
```

```
int mvwgetstr(WINDOW *win,
int y,
```

```
int x,  
char *str);
```

```
int wgetnstr(WINDOW *win,  
char *str,  
int n);
```

```
int wgetstr(WINDOW *win,  
char *str);
```

## Description

The effect of the **getstr** subroutine is as though a series of calls to the **getch** subroutine was made, until a **newline** subroutine, carriage return, or end-of-file is received. The resulting value is placed in the area pointed to by *str*. The string is then terminated with a null byte. The **getnstr**, **mvgetnstr**, **mvwgetnstr**, and **wgetnstr** subroutines read at most *n* bytes, thus preventing a possible overflow of the input buffer. The user's erase and kill characters are interpreted, as well as any special keys (such as function keys, home key, clear key, and so on).

The **mvgetstr** subroutines is identical to the **getstr** subroutine except that it is as though it is a call to the **move** subroutine and then a series of calls to the **getch** subroutine. The **mvwgetstr** subroutine is identical to the **getstr** subroutine except that it is as though it is a call to the **wmove** subroutine and then a series of calls to the **wgetch** subroutine.

The **mvgetnstr** subroutines is identical to the **getstr** subroutine except that it is as though it is a call to the **move** subroutine and then a series of calls to the **getch** subroutine. The **mvwgetnstr** subroutine is identical to the **getstr** subroutine except that it is as though it is a call to the **wmove** subroutine and then a series of calls to the **wgetch** subroutine.

The **getstr**, **wgetstr**, **mvgetstr**, and **mvwgetstr** subroutines will only return the entire multi-byte sequence associated with a character. If the array is large enough to contain at least one character, the subroutines fill the array with complete characters. If the array is not large enough to contain any complete characters, the function fails.

## Parameters

*n*  
*x*  
*y*  
*\*str*      Identifies where to store the string.  
*\*win*      Identifies the window to get the string from and echo it into.

## Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

## Examples

1. To get a string, store it in the user-defined variable *my\_string*, and echo it into the *stdscr*, enter:

```
char *my_string;  
getstr(my_string);
```

2. To get a string, echo it into the user-defined window *my\_window*, and store it in the user-defined variable *my\_string*, enter:

```
WINDOW *my_window;  
char *my_string;  
wgetstr(my_window, my_string);
```

- To get a string in the stdscr at coordinates y=20, x=30, and store it in the user-defined variable `my_string`, enter:
 

```
char *string;
mvgetstr(20, 30, string);
```
- To get a string in the user-defined window `my_window` at coordinates y=20, x=30, and store it in the user-defined variable `my_string`, enter:
 

```
WINDOW *my_window;
char *my_string;
mvwgetstr(my_window, 20, 30, my_string);
```

## Related Information

The **beep** (“beep Subroutine” on page 530) subroutine, **getch** (“getch, mvgetch, mvwgetch, or wgetch Subroutine” on page 561) subroutine, **keypad** (“keypad Subroutine” on page 582) subroutine, **nodelay** (“nodelay Subroutine” on page 599) subroutine, **wgetch** (“getch, mvgetch, mvwgetch, or wgetch Subroutine” on page 561) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## getsyx Subroutine

### Purpose

Retrieves the current coordinates of the virtual screen cursor.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
getsyx(Y, X)
int * Y, * X;
```

### Description

The **getsyx** subroutine retrieves the current coordinates of the virtual screen cursor and stores them in the location specified by `Y` and `X`. The current coordinates are those where the cursor was placed after the last call to the **wnoutrefresh**, **pnoutrefresh**, or **wrefresh**, subroutine. If the **leaveok** subroutine was TRUE for the last window refreshed, then the **getsyx** subroutine returns -1 for both `X` and `Y`.

If lines have been removed from the top of the screen using the **ripoffline** subroutine, `Y` and `X` include these lines. `Y` and `X` should only be used as arguments for the **setsyx** subroutine.

The **getsyx** subroutine, along with the **setsyx** subroutine, is meant to be used by a user-defined function that manipulates curses windows but wants the position of the cursor to remain the same. Such a function would do the following:

- Call the **getsyx** subroutine to obtain the current virtual cursor coordinates.
- Continue manipulating the windows.
- Call the **wnoutrefresh** subroutine on each window manipulated.
- Reset the current virtual cursor coordinates to the original values with the **setsyx** subroutine.
- Refresh the display with a call to the **doupdate** subroutine.

## Parameters

- X* Points to the current row position of the virtual screen cursor. A value of -1 indicates the **leaveok** subroutine was TRUE for the last window refreshed.
- Y* Points to the current column position of the virtual screen cursor. A value of -1 indicates the **leaveok** subroutine was TRUE for the last window refreshed.

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Controlling the Cursor with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## getyx Macro

### Purpose

Returns the coordinates of the logical cursor in the specified window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
getyx( Window, Line, Column)  
WINDOW *Window;  
int Line, Column;
```

### Description

The **getyx** macro returns the coordinates of the logical cursor in the specified window.

## Parameters

- Window* Identifies the window to get the cursor location from.
- Column* Holds the column coordinate of the logical cursor.
- Line* Holds the line or row coordinate of the logical cursor.

## Example

To get the location of the logical cursor in the user-defined window `my_window` and then put these coordinates in the user-defined integer variables `Line` and `Column`, enter:

```
WINDOW *my_window;  
int line, column;  
getyx(my_window, line, column);
```

## Related Information

Controlling the Cursor with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## halfdelay Subroutine

### Purpose

Controls input character delay mode.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include < curses.h >
int halfdelay(int tenths);
```

### Description

The **halfdelay** subroutine sets the input mode for the current window to Half-Delay Mode and specifies tenths of seconds as the half-delay interval. The *tenths* argument must be in a range from 1 up to and including 255.

### Flag

x      Instructs **wgetch** to wait x tenths of a second for input before timing out.

### Parameters

*tenths*

### Return Values

Upon successful completion, the **halfdelay** subroutine returns OK. Otherwise, it returns ERR.

### Related Information

The **cbreak** ("cbreak, nocbreak, noraw, or raw Subroutine" on page 535) subroutine.

---

## has\_colors Subroutine

### Purpose

Determines whether a terminal supports color.

### Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>
has_colors()
```

## Description

The **has\_colors** subroutine determines whether a terminal supports color. If the terminal supports color, the **has\_colors** subroutine returns TRUE. Otherwise, it returns FALSE. Because this subroutine tests for color, you can call it before the **start\_color** subroutine.

The **has\_colors** routine makes writing terminal-independent programs easier because you can use the subroutine to determine whether to use color or another video attribute.

Use the **can\_change\_colors** subroutine to determine whether a terminal that supports colors also supports changing its color definitions.

## Examples

To determine whether or not a terminal supports color, use:

```
has_colors();
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Video Attributes in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## has\_ic and has\_il Subroutine

### Purpose

Query functions for terminal insert and delete capability.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
bool has_ic(void);
bool has_il(void);
```

### Description

The **has\_ic** subroutine indicates whether the terminal has insert- and delete-character capabilities.

The **has\_il** subroutine indicates whether the terminal has insert- and delete-line capabilities, or can simulate them using scrolling regions.

## Return Values

The **has\_ic** subroutine returns a value of TRUE if the terminal has insert- and delete-character capabilities. Otherwise, it returns FALSE.

The **has\_il** subroutine returns a value of TRUE if the terminal has insert- and delete-line capabilities. Otherwise, it returns FALSE.

## Examples

For the **has\_ic** subroutine:

To determine the insert capability of a terminal by returning TRUE or FALSE into the user-defined variable `insert_cap`, enter:

```
int insert_cap;  
insert_cap = has_ic();
```

For the **has\_il** subroutine:

To determine the insert capability of a terminal by returning TRUE or FALSE into the user-defined variable `insert_line`, enter:

```
int insert_line;  
insert_line = has_il();
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## has\_il Subroutine

### Purpose

Determines whether the terminal has insert-line capability.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>  
has_il( )
```

### Description

The **has\_il** subroutine determines whether a terminal has insert-line capability.

### Return Values

The **has\_il** subroutine returns TRUE if terminal has insert-line capability and FALSE, if not.

## Examples

To determine the insert capability of a terminal by returning TRUE or FALSE into the user-defined variable `insert_line`, enter:

```
int insert_line;
insert_line = has_il();
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## idlok Subroutine

### Purpose

Allows curses to use the hardware insert/delete line feature.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
idlok( Window, Flag)
WINDOW *Window;
bool Flag;
```

### Description

The **idlok** subroutine enables curses to use the hardware insert/delete line feature for terminals so equipped. If this feature is disabled, curses cannot use it. The insert/delete line feature is always considered. Enable this option only if your application needs the insert/delete line feature; for example, for a screen editor. If the insert/delete line feature cannot be used, curses will redraw the changed portions of all lines that do not match the desired line.

### Parameters

*Flag*            Specifies whether to enable curses to use the hardware insert/delete line feature (True) or not (False).  
*Window*        Specifies the window it will affect.

## Examples

1. To enable curses to use the hardware insert/delete line feature in `stdscr`, enter:

```
idlok(stdscr, TRUE);
```

2. To force curses not to use the hardware insert/delete line feature in the user-defined window `my_window`, enter:

```
idlok(my_window, FALSE);
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Setting Video Attributes and Curses Options in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## inch, mvinch, mvwinch, or winch Subroutine

### Purpose

Inputs a single-byte character and rendition from a window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
ctype inch(void);
ctype mvinch(int y,
int x);

ctype mvwinch(WINDOW *win,
int y,
int x);
ctype winch(WINDOW *win);
```

### Description

The **inch**, **winch**, **mvinch**, and **mvwinch** subroutines return the character and rendition, of type `ctype`, at the current or specified position in the current or specified window.

### Parameters

*\*win*     Specifies the window from which to get the character.  
*x*  
*y*

### Return Values

Upon successful completion, these subroutines return the specified character and rendition. Otherwise, they return (`ctype`) `ERR`.

### Examples

1. To get the character at the current cursor location in the `stdscr`, enter:

```
ctype character;

character = inch();
```
2. To get the character at the current cursor location in the user-defined window `my_window`, enter:

```
WINDOW *my_window;
chtype character;
```

```
character = winch(my_window);
```

3. To move the cursor to the coordinates  $y = 0$ ,  $x = 5$  and then get that character, enter:

```
chtype character;
```

```
character = mvinch(0, 5);
```

4. To move the cursor to the coordinates  $y = 0$ ,  $x = 5$  in the user-defined window `my_window` and then get that character, enter:

```
WINDOW *my_window;
chtype character;
```

```
character = mvwinch(my_window, 0, 5);
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## init\_color Subroutine

### Purpose

Changes a color definition.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
init_color( Color, R,  
G, B)
```

```
register short Color, R, G, B;
```

### Description

The **init\_color** subroutine changes a color definition. A single color is defined by the combination of its red, green, and blue components. The **init\_color** subroutine changes all the occurrences of the color on the screen immediately. If the color is changed successfully, this subroutine returns OK. Otherwise, it returns ERR.

**Note:** The values for the red, green, and blue components must be between 0 (no component) and 1000 (maximum amount of component). The **init\_color** subroutine sets values less than 0 to 0 and values greater than 1000 to 1000.

To determine if you can change a terminal's color definitions, see the **can\_change\_color** subroutine.

## Return Values

- OK** Indicates the color was changed successfully.  
**ERR** Indicates the color was not changed.

## Parameters

- Color* Identifies the color to change. The value of the parameter must be between **0** and **COLORS-1**.  
*R* Specifies the desired intensity of the red component.  
*G* Specifies the desired intensity of the green component.  
*B* Specifies the desired intensity of the blue component.

## Examples

To initialize the color definition for color 11 to violet on a terminal that supports at least 12 colors, use:  
`init_color(11,500,0,500);`

## Related Information

The **start\_color** (“start\_color Subroutine” on page 636) subroutine.

Curses Overview for Programming and Manipulating Video Attributes in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## init\_pair Subroutine

### Purpose

Changes a color-pair definition.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
init_pair( Pair, F, B)  
register short Pair, F, B;
```

### Description

The **init\_pair** subroutine changes a color-pair definition. A color pair is a combination of a foreground and a background color. If you specify a color pair that was previously initialized, curses refreshes the screen and changes all occurrences of that color pair to the new definition. You must call the **start\_color** subroutine before you call this subroutine.

## Return Values

- OK** Indicates successful completion.  
**ERR** Indicates the subroutine failed.

## Parameters

- Pair* Identifies the color-pair number. The value of the *Pair* parameter must be between **1** and **COLORS\_PAIRS-1**.
- F* Specifies the foreground color number. This number must be between **0** and **COLORS-1**.
- B* Specifies the background color number. This number must be between **0** and **COLORS-1**.

## Examples

To initialize the color definition for color-pair 2 to a black foreground (color 0) with a cyan background (color 3), use:

```
init_pair(2,COLOR_BLACK, COLOR_CYAN);
```

## Related Information

The **init\_color** (“init\_color Subroutine” on page 575) subroutine, **start\_color** (“start\_color Subroutine” on page 636) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Video Attributes in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## initscr and newterm Subroutine

### Purpose

Initializes curses and its data structures.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
WINDOW *initscr(void);
SCREEN *newterm(char *type,
FILE *outfile,
FILE *infile);
```

### Description

The **initscr** subroutine determines the terminal type and initializes all implementation data structures. The TERM environment variable specifies the terminal type. The **initscr** subroutine also causes the first refresh operation to clear the screen. If errors occur, **initscr** writes an appropriate error message to standard error and exits. The only subroutines that can be called before **initscr** or **newterm** are the **filter**, **ripoffline**, **slk\_init**, **use\_env**, and the subroutines whose prototypes are defined in <term.h>. Portable applications must not call **initscr** twice.

The **newterm** subroutine can be called as many times as desired to attach a terminal device. The type argument points to a string specifying the terminal type, except that, if type is a null pointer, the TERM environment variable is used. The *outfile* and *infile* arguments are file pointers for output to the terminal and input from the terminal, respectively. It is unspecified whether Curses modifies the buffering mode of these file pointers. The **newterm** subroutine should be called once for each terminal.

The **initscr** subroutine is equivalent to:

```
newterm(gentenv("TERM"), stdout, stdin); return stdscr;
```

If the current disposition for the signals SIGINT, SIGQUIT or SIGTSTP is SIGDFL, then the **initscr** subroutine may also install a handler for the signal, which may remain in effect for the life of the process or until the process changes the disposition of the signal.

The **initscr** and **newterm** subroutines initialise the `cur_term` external variable.

**initscr**                    CURSES                    Curses Interfaces

## Return Values

Upon successful completion, the **initscr** subroutine returns a pointer to **stdscr**. Otherwise, it does not return.

Upon successful completion, the **newterm** subroutine returns a pointer to the specified terminal. Otherwise, it returns a null pointer.

## Example

To initialize curses so that other curses subroutines can be called, use:

```
initscr();
```

## Related Information

The **doupdate** (“doupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine, **del\_curterm** (“del\_curterm, restartterm, set\_curterm, or setupterm Subroutine” on page 547) subroutine, **filter** (“filter Subroutine” on page 557) subroutine, **slk\_attroff** (“slk\_attroff, slk\_attr\_off, slk\_attron, slk\_attrset, slk\_attr\_set, slk\_clear, slk\_color, slk\_init, slk\_label, slk\_noutrefresh, slk\_refresh, slk\_restore, slk\_set, slk\_touch, slk\_wset, Subroutine” on page 627) subroutine, **setupterm** (“setupterm Subroutine” on page 624) subroutine.

Curses Overview for Programming, Initializing Curses, List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## insch, mvinsch, mvwinsch, or winsch Subroutine

### Purpose

Inserts a single-byte character and rendition in a window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
int insch(chtype ch);
int mvinsch(int y,
chtype h);

int mvwinsch(WINDOW *win,
int x,
int y,
chtype h);
int winsch(WINDOW *win,
chtype h);
```

## Description

These subroutines insert the character and rendition into the current or specified window at the current or specified position.

These subroutines do not perform wrapping or advance the cursor position. These functions perform special-character processing, with the exception that if a **newline** is inserted into the last line of a window and scrolling is not enabled, the behavior is unspecified.

## Parameters

*ch*  
*y*  
*x*  
*\*win* Specifies the window in which to insert the character.

## Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

## Examples

1. To insert the character *x* in the stdscr, enter:

```
ctype x;  
insch(x);
```

2. To insert the character *x* into the user-defined window *my\_window*, enter:

```
WINDOW *my_window  
ctype x;  
winsch(my_window, x);
```

3. To move the logical cursor to the coordinates *Y=10, X=5* prior to inserting the character *x* in the stdscr, enter:

```
ctype x;  
mvinsch(10, 5, x);
```

4. To move the logical cursor to the coordinates *y=10, X=5* prior to inserting the character *x* in the user-defined window *my\_window*, enter:

```
WINDOW *my_window;  
ctype x;  
mwinsch(my_window, 10, 5, x);
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## insertln or winsertln Subroutine

### Purpose

Inserts a blank line above the current line in a window.

## Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>
int insertln(void)
```

```
int winsertln(WINDOW *win);
```

## Description

The **insertln** and **winsertln** subroutines insert a blank line before the current line in the current or specified window. The bottom line is no longer displayed. The cursor position does not change.

## Parameters

*\*win* Specifies the window in which to insert the blank line.

## Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

## Examples

1. To insert a blank line above the current line in the stdscr, enter:

```
insertln();
```

2. To insert a blank line above the current line in the user-defined window `my_window`, enter:

```
WINDOW *mywindow;
winsertln(my_window);
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## intrflush Subroutine

### Purpose

Enables or disables flush on interrupt.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int intrflush(WINDOW * win,
bool bf);
```

## Description

The **intrflush** subroutine specifies whether pressing an interrupt key (interrupt, suspend, or quit) will flush the input buffer associated with the current screen. If the value of *bf* is TRUE, then flushing of the output buffer associated with the current screen will occur when an interrupt key (interrupt, suspend, or quit) is pressed. If the value of *bf* is FALSE then no flushing of the buffer will occur when an interrupt key is pressed. The default for the option is inherited from the display driver settings. The *win* argument is ignored.

## Parameters

*bf*

*\*win* Specifies the window for which to enable or disable queue flushing.

## Return Values

Upon successful completion, the **intrflush** subroutine returns OK. Otherwise, it returns ERR.

## Examples

1. To enable queue flushing in the user-defined window *my\_window*, enter:  

```
intrflush(my_window, TRUE);
```
2. To disable queue flushing in the user-defined window *my\_window*, enter:  

```
intrflush(my_window, FALSE);
```

## Related Information

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Setting Video Attributes and Curses Options in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## keyname, key\_name Subroutine

### Purpose

Gets the name of keys.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
char *keyname(int c);
```

```
char *key_name(wchar_t c);
```

## Description

The **keyname** and **key\_name** subroutines generate a character string whose value describes the key *c*. The *c* argument of **keyname** can be an 8-bit character or a key code. The *c* argument of **key\_name** must be a wide character.

The string has a format according to the first applicable row in the following table:

Input	Format of Returned String
Visible character	The same character
Control character	^X
Meta-character ( <b>keyname</b> only)	M-X
Key value defined in <curses.h> ( <b>keyname</b> only)	KEY_name
None of the above	UNKNOWN KEY

The meta-character notation shown above is used only, if meta-characters are enabled.

## Parameter

*c*

## Return Values

Upon successful completion, the **keyname** subroutine returns a pointer to a string as described above. Otherwise, it returns a null pointer.

## Examples

```
int key;
char *name;
 keypad(stdscr, TRUE);
 addstr("Hit a key");
 key=getch();
 name=keyname(key);
```

**Note:** If the Page Up key is pressed, **keyname** will return **KEY\_PPAGE**.

## Related Information

The **meta** (“meta Subroutine” on page 588) and **wgetch** (“getch, mvgetch, mvwgetch, or wgetch Subroutine” on page 561) subroutines.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## keypad Subroutine

### Purpose

Enables or disables abbreviation of function keys.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int keypad(WINDOW *win,  
bool bf);
```

## Description

The **keypad** subroutine controls keypad translation. If *bf* is TRUE, keypad translation is turned on. If *bf* is FALSE, keypad translation is turned off. The initial state is FALSE.

This subroutine affects the behavior of any function that provides keyboard input.

If the terminal in use requires a command to enable it to transmit distinctive codes when a function key is pressed, then after keypad translation is first enabled, the implementation transmits this command to the terminal before an affected input function tries to read any characters from that terminal.

## Parameters

*bf*  
*\*win* Specifies the window in which to enable or disable the keypad.

## Return Values

Upon successful completion, the **keypad** subroutine returns OK. Otherwise, it returns ERR.

## Examples

To turn on the keypad in the user-defined window *my\_window*, use:

```
WINDOW *my_window;  
keypad(my_window, TRUE);
```

## Related Information

The **getch** (“getch, mvgetch, mvwgetch, or wgetch Subroutine” on page 561) subroutine.

The **terminfo** file format.

Curses Overview for Programming, List of Curses Subroutines, Setting Video Attributes and Curses Options in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## killchar or killwchar Subroutine

### Purpose

Terminal environment query functions.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>  
char killchar(void);  
int killwchar(wchar_t *ch);
```

### Description

The **killchar** subroutine returns the current line.

The **killchar** subroutine stores the current line kill character in the object pointed to by *ch*. If no line kill character has been defined, the subroutine will fail and the object pointed to by *ch* will not be changed.

## Parameters

*\*ch*

## Return Values

The **killchar** subroutine returns the line kill character. The return value is unspecified when this character is a multi-byte character.

Upon successful completion, the **killchar** subroutine returns OK. Otherwise, it returns ERR.

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## **\_lazySetErrorHandler** Subroutine

### Purpose

Installs an error handler into the lazy loading runtime system for the current process.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <sys/ldr.h>
#include <sys/errno.h>
typedef void *handler_t
char *module;
char *symbol;
unsigned int errval;
handler_t *_lazySetErrorHandler
handler_t *err_handler;
```

### Description

This function allows a process to install a custom error handler to be called when a lazy loading reference fails to find the required module or function. This function should only be used when the main program or one of its dependent modules was linked with the **-blazy** option. To call **\_lazySetErrorHandler** from a module that is not linked with the **-blazy** option, you must use the **-lrtl** option. If you use **-blazy**, you do not need to specify **-lrtl**.

This function is not thread safe. The calling program should ensure that **\_lazySetErrorHandler** is not called by multiple threads at the same time.

The user-supplied error handler may print its own error message, provide a substitute function to be used in place of the called function, or call **longjmp** subroutine. To provide a substitute function that will be called instead of the originally referenced function, the error handler should return a pointer

## Parameters

<i>Column</i>	Specifies the horizontal position to move the logical cursor to before getting the character.
<i>Line</i>	Specifies the vertical position to move the logical cursor to before getting the character.
<i>Window</i>	Identifies the window to get the character from and echo it into.

## Return Values

Upon completion, the character code for the data key or one of the following values is returned:

<b>KEY_</b> <i>xxxx</i>	The <b>keypad</b> subroutine is set to TRUE and a control key was recognized. See the <b>curses.h</b> file for a complete list of the key codes that can be returned.
-------------------------	---

## Examples

1. To get a character and echo it to the stdscr, use:

```
mvgetch();
```

2. To get a character and echo it into stdscr at the coordinates y=20, x=30, use:

```
mvgetch(20, 30);
```

3. To get a character and echo it into the user-defined window `my_window` at coordinates y=20, x=30, use:

```
WINDOW *my_window;  
mvwgetch(my_window, 20, 30);
```

## Related Information

The **keypad** (“keypad Subroutine” on page 582) subroutine, **meta** (“meta Subroutine” on page 588) subroutine, **nodelay** (“nodelay Subroutine” on page 599) subroutine, **echo** or **noecho** (“echo or noecho Subroutine” on page 552) subroutine, **notimeout** (“notimeout, timeout, wtimeout Subroutine” on page 599) subroutine, **ebreak** or **nocbreak** (“cbreak, nocbreak, noraw, or raw Subroutine” on page 535) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## leaveok Subroutine

### Purpose

Controls physical cursor placement after a call to the **refresh** subroutine.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
leaveok( Window, Flag)  
WINDOW *Window;  
bool Flag;
```

## Description

The **leaveok** subroutine controls cursor placement after a call to the **refresh** (“refresh or wrefresh Subroutine” on page 608) subroutine. If the *Flag* parameter is set to **FALSE**, curses leaves the physical cursor in the same location as logical cursor when the window is refreshed.

If the *Flag* parameter is set to **TRUE**, curses leaves the cursor as is and does not move the physical cursor when the window is refreshed. This option is useful for applications that do not use the cursor, because it reduces physical cursor motions.

By default **leaveok** is **FALSE**, and the physical cursor is moved to the same position as the logical cursor after a refresh.

## Parameters

*Flag* Specifies whether to leave the physical cursor alone after a refresh (**TRUE**) or to move the physical cursor to the logical cursor after a refresh (**FALSE**).

*Window* Identifies the window to set the *Flag* parameter for.

## Return Values

**OK** Indicates the subroutine completed. The **leaveok** subroutine always returns this value.

## Examples

1. To move the physical cursor to the same location as the logical cursor after refreshing the user-defined window *my\_window*, enter:

```
WINDOW *my_window;  
leaveok(my_window, FALSE);
```

2. To leave the physical cursor alone after refreshing the user-defined window *my\_window*, enter:

```
WINDOW *my_window;  
leaveok(my_window, TRUE);
```

## Related Information

The **refresh** (“refresh or wrefresh Subroutine” on page 608) subroutine.

Controlling the Cursor with Curses, Curses Overview for Programming, List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## longname Subroutine

### Purpose

Returns the verbose name of a terminal.

### Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>
char *longname(void);
```

## Description

The **longname** subroutine generates a verbose description for the current terminal. The maximum length of a verbose description is 128 bytes. It is defined only after the call to the **initscr** or **newterm** subroutines.

The area is overwritten by each call to the **newterm** subroutine, so the value should be saved if you plan on using the **longname** subroutine with multiple terminals.

## Return Values

Upon successful completion, the **longname** subroutine returns a pointer to the description specified above. Otherwise, it returns a null pointer on error.

## Related Information

The **initscr** (“initscr and newterm Subroutine” on page 577) subroutine, **newterm** (“newterm Subroutine” on page 594) subroutine, **setupterm** (“setupterm Subroutine” on page 624) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## makenew Subroutine

### Purpose

Creates a new window buffer and returns a pointer.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
WINDOW *makenew( )
```

### Description

The **makenew** subroutine creates a new window buffer and returns a pointer to it. The **makenew** subroutine is called by the **newwin** subroutine to create the window structure. The **makenew** subroutine should not be called directly by a program.

### Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Windows with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## meta Subroutine

### Purpose

Enables/disables meta-keys.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>

int meta(WINDOW *win,
bool bf);
```

### Description

Initially, whether the terminal returns 7 or 8 significant bits on input depends on the control mode of the display driver. To force 8 bits to be returned, invoke the **meta** subroutine (win, TRUE). To force 7 bits to be returned, invoke the **meta** subroutine (win, FALSE). The *win* argument is always ignored.

If the terminfo capabilities **smm** (meta\_on) and **rmm** (meta\_off) are defined for the terminal, **smm** is sent to the terminal when **meta** (win, TRUE) is called and **rmm** is sent when **meta** (win, FALSE) is called.

### Parameters

*bf*  
*\*win*

### Return Values

Upon successful completion, the **meta** subroutine returns OK. Otherwise, it returns ERR.

### Examples

1. To request an 8-bit character return when using a **getch** routine, enter:

```
WINDOW *some_window;
meta(some_window, TRUE);
```

2. To strip the highest bit off the character returns in the user-defined window *my\_window*, enter:

```
WINDOW *some_window;
meta(some_window, FALSE);
```

### Related Information

The **getch** (“getch, mvgetch, mvwgetch, or wgetch Subroutine” on page 561) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## move or wmove Subroutine

### Purpose

Window location cursor functions.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
int ( x);
```

```
int wmove (WINDOW *win,  
int y,  
int x);
```

### Description

The **move** and **wmove** subroutines move the logical cursor associated with the current or specified window to (y, x) relative to the window's origin. This subroutine does not move the cursor of the terminal until the next **refresh** ("refresh or wrefresh Subroutine" on page 608) operation.

### Parameters

y  
x  
\*win

### Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

### Examples

1. To move the logical cursor in the stdscr to the coordinates y = 5, x = 10, use:  

```
move(5, 10);
```
2. To move the logical cursor in the user-defined window `my_window` to the coordinates y = 5, x = 10, use:  

```
WINDOW *my_window;  
wmove(my_window, 5, 10);
```

### Related Information

The **getch** ("getch, mvgetch, mvwgetch, or wgetch Subroutine" on page 561) and **refresh** ("refresh or wrefresh Subroutine" on page 608) subroutines.

Controlling the Cursor with Curses, Curses Overview for Programming, List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## mvcur Subroutine

### Purpose

Output cursor movement commands to the terminal.

## Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>

int mvcur(int oldrow,
int oldcol,
int newrow,
int newcol);
```

## Description

The **mvcur** subroutine outputs one or more commands to the terminal that move the terminal's cursor to (*newrow*, *newcol*), an absolute position on the terminal screen. The (*oldrow*, *oldcol*) arguments specify the former cursor position. Specifying the former position is necessary on terminals that do not provide coordinate-based movement commands. On terminals that provide these commands, Curses may select a more efficient way to move the cursor based on the former position. If (*newrow*, *newcol*) is not a valid address for the terminal in use, the **mvcur** subroutine fails. If (*oldrow*, *oldcol*) is the same as (*newrow*, *newcol*), **mvcur** succeeds without taking any action. If **mvcur** outputs a cursor movement command, it updates its information concerning the location of the cursor on the terminal.

## Parameters

*newcol*  
*newrow*  
*oldcol*  
*oldrow*

## Return Values

Upon successful completion, the **mvcur** subroutine returns OK. Otherwise, it returns ERR.

## Examples

1. To move the physical cursor from the coordinates  $y = 5$ ,  $x = 15$  to  $y = 25$ ,  $x = 30$ , use:  

```
mvcur(5, 15, 25, 30);
```
2. To move the physical cursor from unknown coordinates to  $y = 5$ ,  $x = 0$ , use:  

```
mvcur(50, 50, 5, 0);
```

In this example, the physical cursor's current coordinates are unknown. Therefore, arbitrary values are assigned to the *OldLine* and *OldColumn* parameters and the desired coordinates are assigned to the *NewLine* and *NewColumn* parameters. This is called an *absolute move*.

## Related Information

The **doupdate** ("doupdate, refresh, wnoutrefresh, or wrefresh Subroutines" on page 657) subroutine, **is\_linetouched** ("is\_linetouched, is\_wintouched, touchline, touchwin, untouchwin, or wtouchin Subroutine" on page 648) subroutine, **move** ("move or wmove Subroutine" on page 589) subroutine, **refresh** ("refresh or wrefresh Subroutine" on page 608) subroutine.

Controlling the Cursor with Curses, Curses Overview for Programming, List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## mvwin Subroutine

### Purpose

Moves a window or subwindow to the specified coordinates.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int mvwin  
(WINDOW *win,  
int y,  
int x);
```

### Description

The **mvwin** subroutine moves the specified window so that its origin is at position (y, x). If the move causes any portion of the window to extend past any edge of the screen, the function fails and the window is not moved.

### Parameters

\*win  
x  
y

### Return Values

Upon successful completion, the **mvwin** subroutine returns OK. Otherwise, it returns ERR.

### Examples

1. To move the user-defined window `my_window` from its present location to the upper left corner of the terminal, enter:

```
WINDOW *my_window;  
mvwin(my_window, 0, 0);
```

2. To move the user-defined window `my_window` from its present location to the coordinates y = 20, x = 10, enter:

```
WINDOW *my_window;  
mvwin(my_window, 20, 10);
```

### Related Information

The **derwin** (“derwin, newwin, or subwin Subroutine” on page 596) subroutine, **doupdate** (“doupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine, **is\_linetouched** (“is\_linetouched, is\_wintouched, touchline, touchwin, untouchwin, or wtouchin Subroutine” on page 648) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

## newpad, pnoutrefresh, prefresh, or subpad Subroutine

### Purpose

Pad management functions.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
WINDOW *newpad  
(int nlines,  
int ncols);
```

```
int  
pnoutrefresh  
(WINDOW *pad,  
int pminrow,  
int pmincol,  
int sminrow,  
int smincol,  
int smaxrow,  
int smaxcol);
```

```
int  
prefresh  
(WINDOW *pad,  
int pminrow,  
int pmincol,  
int sminrow,  
int smincol,  
int smaxrow,  
int smaxcol);
```

```
WINDOW  
*subpad  
(WINDOW *orig,  
int nlines,  
int ncols,  
int begin_y,  
int begin_x);
```

### Description

The **newpad** subroutine creates a specialised WINDOW data structure with *nlines* lines and *ncols* columns. A pad is similar to a window, except that it is not associated with a viewable part of the screen. Automatic refreshes of pads do not occur.

The **subpad** subroutine creates a subwindow within a pad with *nlines* lines and *ncols* columns. Unlike the **subwin** subroutine, which uses screen coordinates, the window is at a position (*begin\_y*, *begin\_x*) on the pad. The window is made in the middle of the window *orig*, so that changes made to one window affects both windows.

The **prefresh** (“prefresh or pnoutrefresh Subroutine” on page 603) or **pnoutrefresh** (“prefresh or pnoutrefresh Subroutine” on page 603) subroutines are analogous to the **wrefresh** and **wnoutrefresh** subroutines except that they relate to pads instead of windows. The additional arguments indicate what

part of the pad and screen are involved. The *pminrow* and *pmincol* arguments specify the origin of the rectangle to be displayed in the screen. The lower right-hand corner of the rectangle to be displayed in the pad is calculated from the screen coordinates, since the rectangles must be the same size. Both rectangles must be entirely contained within their respective structures. Negative values of *pminrow*, *pmincol*, *sminrow* or *smincol* are treated as if they were zero.

## Parameters

*ncols*  
*nlines*  
*begin\_x*  
*begin\_y*  
*\*orig*  
*\*pad*  
*pminrow*  
*pmincol*  
*sminrow*  
*smincol*  
*smaxrow*  
*smaxcol*

## Return Values

Upon successful completion, the **newpad** and **subpad** subroutines return a pointer to the pad structure. Otherwise, they return a null pointer.

Upon successful completion, the **pnoutrefresh** and **prefresh** subroutines return OK. Otherwise, they return ERR.

## Examples

For the **newpad** subroutine:

1. To create a new pad and save the pointer to it in `my_pad`, enter:

```
WINDOW *my_pad;  
  
my_pad = newpad(5, 10);
```

`my_pad` is now a pad 5 lines deep, 10 columns wide.

2. To create a pad and save the pointer to it in `my_pad`, which is flush with the right side of the terminal, enter:

```
WINDOW *my_pad;  
  
my_pad = newpad(5, 0);
```

`my_pad` is now a pad 5 lines deep, extending to the far right side of the terminal.

3. To create a pad and save the pointer to it in `my_pad`, which fills the entire terminal, enter:

```
WINDOW *my_pad;  
  
my_pad = newpad(0, 0);
```

`my_pad` is now a pad that fills the entire terminal.

4. To create a very large pad and display part of it on the screen, enter;

```
WINDOW *my_pad;

my_pa1 = newpad(120,120);

prefresh (my_pa1, 0,0,0,0,20,30);
```

This causes the first 21 rows and first 31 columns of the pad to be displayed on the screen. The upper left coordinates of the resulting rectangle are (0,0) and the bottom right coordinates are (20,30).

For the **prefresh** or **pnoutrefresh** subroutines:

1. To update the user-defined my\_pad pad from the upper-left corner of the pad on the terminal with the upper-left corner at the coordinates Y=20, X=10 and the lower-right corner at the coordinates Y=30, X=25 enter

```
WINDOW *my_pad;
prefresh(my_pad, 0, 0, 20, 10, 30, 25);
```

2. To update the user-defined my\_pad1 and my\_pad2 pads and output them both to the terminal in one burst of output, enter:

```
WINDOW *my_pad1; *my_pad2;
pnoutrefresh(my_pad1, 0, 0, 20, 10, 30, 25);
pnoutrefresh(my_pad2, 0, 0, 0, 0, 10, 5);
doupdate();
```

For the **subpad** subroutine:

To create a subpad, use:

```
WINDOW *orig, *mypad;
orig = newpad(100, 200);
mypad = subpad(orig, 30, 5, 25, 180);
```

The parent pad is 100 lines by 200 columns. The subpad is 30 lines by 5 columns and starts in line 25, column 180 of the parent pad.

## Related Information

The **derwin** (“derwin, newwin, or subwin Subroutine” on page 596) subroutine, **doupdate** (“doupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine, **is\_linetouched** (“is\_linetouched, is\_wintouched, touchline, touchwin, untouchwin, or wtouchin Subroutine” on page 648) subroutine, **prefresh** (“prefresh or pnoutrefresh Subroutine” on page 603) or **pnoutrefresh** (“prefresh or pnoutrefresh Subroutine” on page 603) subroutine, and **subpad** (“subpad Subroutine” on page 637) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Windows in the Curses Environment in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## newterm Subroutine

### Purpose

Initializes curses and its data structures for a specified terminal.

### Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>
```

```
SCREEN *newterm(  
    Type,  
    OutFile, InFile)  
char *Type;  
FILE *OutFile, *InFile;
```

## Description

The **newterm** subroutine initializes **curses** and its data structures for a specified terminal. Use this subroutine instead of the **initscr** subroutine if you are writing a program that sends output to more than one terminal. You should also use this subroutine if your program requires indication of error conditions so that it can run in a line-oriented mode on terminals that do not support a screen-oriented program.

If you are directing your program's output to more than one terminal, you must call the **newterm** subroutine once for each terminal. You must also call the **endwin** subroutine for each terminal to stop **curses** and restore the terminal to its previous state.

## Parameters

<i>InFile</i>	Identifies the input device file.
<i>OutFile</i>	Identifies the output device file.
<i>Type</i>	Specifies the type of output terminal. This parameter is the same as the <b>\$TERM</b> environment variable for that terminal.

## Return Values

The **newterm** subroutine returns a variable of type **SCREEN \***. You should save this reference to the terminal within your program.

## Examples

1. To initialize **curses** on a terminal represented by the `lft` device file as both the input and output terminal, open the device file with the following:

```
fdfile = fopen("/dev/lft0", "r+");
```

Then, use the **newterm** subroutine to initialize **curses** on the terminal and save the new terminal in the *my\_terminal* variable as follows:

```
char termname [] = "terminaltype";  
SCREEN *my_terminal;  
my_terminal = newterm(termname,fdfile, fdfile);
```

2. To open the device file `/dev/lft0` as the input terminal and the `/dev/tty0` (an `ibm3151`) as the output terminal, do the following:

```
fdifile = fopen("/dev/lft0", "r");  
fdofile = fopen("/dev/tty0", "w");
```

```
SCREEN *my_terminal2;  
my_terminal2 = newterm("ibm3151",fdofile, fdifile);
```

3. To use `stdin` for input and `stdout` for output, do the following:

```
char termname [] = "terminaltype";  
SCREEN *my_terminal;  
my_terminal = newterm(termname,stdout,stdin);
```

## Related Information

The **endwin** (“endwin Subroutine” on page 554) subroutine, **initscr** (“initscr and newterm Subroutine” on page 577) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Initializing Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## derwin, newwin, or subwin Subroutine

### Purpose

Window creation subroutines.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
WINDOW *derwin(WINDOW *orig,  
int nlines,  
int ncols,  
int begin_y,  
int begin_x);
```

```
WINDOW *newwin(int nlines,  
int ncols,  
int begin_y,  
int begin_x);
```

```
WINDOW *subwin(WINDOW *orig,  
int nlines,  
int ncols,  
int begin_y,  
int begin_x);
```

### Description

The **derwin** subroutine is the same as the **subwin** subroutine except that *begin\_y* and *begin\_x* are relative to the origin of the window *orig* rather than absolute screen positions.

The **newwin** subroutine creates a new window with *nlines* lines and *ncols* columns, positioned so that the origin is at (*begin\_y*, *begin\_x*). If *nlines* is zero, it defaults to `LINES - begin_y`; if *ncols* is zero, it defaults to `COLS - begin_x`.

The **subwin** subroutine creates a new window with *nlines* lines and *ncols* columns, positioned so that the origin is at (*begin\_y*, *begin\_x*). (This position is an absolute screen position, not a position relative to the window origin.) If any part of the new window is outside *orig*, the subroutine fails and the window is not created.

### Parameters

*ncols*  
*nlines*  
*begin\_y*  
*begin\_x*

## Return Values

Upon successful completion, these subroutines return a pointer to the new window. Otherwise, they return a null pointer.

## Examples

For the **derwin** and **newwin** subroutines:

1. To create a new window, enter:

```
WINDOW *my_window;
```

```
my_window = newwin(5, 10, 20, 30);
```

`my_window` is now a window 5 lines deep, 10 columns wide, starting at the coordinates `y = 20, x = 30`. That is, the upper left corner is at coordinates `y = 20, x = 30`, and the lower right corner is at coordinates `y = 24, x = 39`.

2. To create a window that is flush with the right side of the terminal, enter:

```
WINDOW *my_window;
```

```
my_window = newwin(5, 0, 20, 30);
```

`my_window` is now a window 5 lines deep, extending all the way to the right side of the terminal, starting at the coordinates `y = 20, x = 30`. The upper left corner is at coordinates `y = 20, x = 30`, and the lower right corner is at coordinates `y = 24, x = lastcolumn`.

3. To create a window that fills the entire terminal, enter:

```
WINDOW *my_window;
```

```
my_window = newwin(0, 0, 0, 0);
```

`my_window` is now a screen that is a window that fills the entire terminal's display.

For the **subwin** subroutine:

1. To create a subwindow, use:

```
WINDOW *my_window, *my_sub_window;
```

```
my_window = newwin ("derwin, newwin, or subwin Subroutine" on page 596)  
                (5, 10, 20, 30);
```

`my_sub_window` is now a subwindow 2 lines deep, 5 columns wide, starting at the same coordinates of its parent window `my_window`. That is, the subwindow's upper-left corner is at coordinates `y = 20, x = 30` and lower-right corner is at coordinates `y = 21, x = 34`.

2. To create a subwindow that is flush with the right side of its parent, use

```
WINDOW *my_window, *my_sub_window;
```

```
my_window =  
newwin ("derwin, newwin, or subwin Subroutine" on page 596)(5, 10, 20, 30);  
my_sub_window = subwin(my_window, 2, 0, 20, 30);
```

`my_sub_window` is now a subwindow 2 lines deep, extending all the way to the right side of its parent window `my_window`, and starting at the same coordinates. That is, the subwindow's upper-left corner is at coordinates `y = 20, x = 30` and lower-right corner is at coordinates `y = 21, x = 39`.

3. To create a subwindow in the lower-right corner of its parent, use:

```
WINDOW *my_window, *my_sub_window
```

```
my_window = newwin ("derwin, newwin, or subwin Subroutine" on page 596)
```

```
(5, 10, 20, 30);  
my_sub_window = subwin(my_window, 0, 0, 22, 35);
```

`my_sub_window` is now a subwindow that fills the bottom right corner of its parent window, `my_window`, starting at the coordinates `y = 22, x = 35`. That is, the subwindow's upper-left corner is at coordinates `y = 22, x = 35` and lower-right corner is at coordinates `y = 24, x = 39`.

## Related Information

The **endwin** (“endwin Subroutine” on page 554), **initscr** (“initscr and newterm Subroutine” on page 577) subroutines.

Curses Overview for Programming, List of Curses Subroutines, Windows in the Curses Environment in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## nl or nonl Subroutine

### Purpose

Enables/disables newline translation.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>  
int nl(void);  
int nonl(void);
```

### Description

The **nl** subroutine enables a mode in which carriage return is translated to newline on input. The **nonl** subroutine disables the above translation. Initially, the above translation is enabled.

### Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

### Examples

1. To instruct **wgetch** to translate the carriage return into a newline, enter:  
`nl();`
2. To instruct **wgetch** not to translate the carriage return, enter:  
`nonl();`

## Related Information

The **refresh** (“refresh or wrefresh Subroutine” on page 608) subroutine, **waddch** (“addch, mvaddch, mvwaddch, or waddch Subroutine” on page 523) subroutine.

Curses Overview for Programming, Understanding Terminals with Curses, List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## nodelay Subroutine

### Purpose

Enables or disables block during read.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>

int nodelay(WINDOW *win,
            bool bf);
```

### Description

The **nodelay** subroutine specifies whether Delay Mode or No Delay Mode is in effect for the screen associated with the specified window. If *bf* is TRUE, this screen is set to No Delay Mode. If *bf* is FALSE, this screen is set to Delay Mode. The initial state is FALSE.

### Parameters

*bf*  
*\*win*

### Return Values

Upon successful completion, the **nodelay** subroutine returns OK. Otherwise, it returns ERR.

### Examples

1. To cause the **wgetch** subroutine to return an error message, if no input is ready in the user-defined window `my_window`, use:  

```
nodelay(my_window, TRUE);
```
2. To allow for a delay when retrieving a character in the user-defined window `my_window`, use:  

```
WINDOW *my_window;
nodelay(my_window, FALSE);
```

### Related Information

The **halfdelay** (“halfdelay Subroutine” on page 570) subroutine, **wgetch** (“getch, mvgetch, mvwgetch, or wgetch Subroutine” on page 561) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*

---

## notimeout, timeout, wtimeout Subroutine

### Purpose

Controls blocking on input.

### Library

Curses Library (**libcurses.a**)

## Curses Syntax

```
#include <curses.h>
```

```
int notimeout  
(WINDOW *win,  
bool bf);
```

```
void timeout  
(int delay);
```

```
void wtimeout  
(WINDOW *win,  
int delay);
```

## Description

The **notimeout** subroutine specifies whether Timeout Mode or No Timeout Mode is in effect for the screen associated with the specified window. If *bf* is TRUE, this screen is set to No Timeout Mode. If *bf* is FALSE, this screen is set to Timeout Mode. The initial state is FALSE.

The **timeout** and **wtimeout** subroutines set blocking or non-blocking read for the current or specified window based on the value of *delay*:

<i>delay</i> < 0	One or more blocking reads (indefinite waits for input) are used.
<i>delay</i> = 0	One or more non-blocking reads are used. Any Curses input subroutine will fail if every character of the requested string is not immediately available.
<i>delay</i> > 0	Any Curses input subroutine blocks for <i>delay</i> milliseconds and fails if there is still no input.

## Parameters

*\*win*  
*bf*

## Return Values

Upon successful completion, the **notimeout** subroutine returns OK. Otherwise, it returns ERR.

The **timeout** and **wtimeout** subroutines do not return a value.

## Examples

To set the flag so that the **wgetch** subroutine does not set the timer when getting characters from the *my\_win* window, use:

```
WINDOW *my_win;  
notimeout(my_win, TRUE);
```

## Related Information

The **getch** (“getch, mvgetch, mvwgetch, or wgetch Subroutine” on page 561), **halfdelay** (“halfdelay Subroutine” on page 570), **nodelay** (“nodelay Subroutine” on page 599), and **notimeout** (“notimeout, timeout, wtimeout Subroutine” on page 599) subroutines.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Getting Characters in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## overlay or overwrite Subroutine

### Purpose

Copies one window on top of another.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
WINDOW *dstwin);  
int overwrite(const WINDOW *srcwin,  
WINDOW *dstwin);
```

### Description

The **overlay** and **overwrite** subroutines overlay *srcwin* on top of *dstwin*. The *srcwin* and *dstwin* arguments need not be the same size; only text where the two windows overlap is copied.

The **overwrite** subroutine copies characters as though a sequence of **win\_wch** and **wadd\_wch** subroutines were performed with the destination window's attributes and background attributes cleared.

The **overlay** subroutine does the same thing, except that, whenever a character to be copied is the background character of the source window, the **overlay** subroutine does not copy the character but merely moves the destination cursor the width of the source background character.

If any portion of the overlaying window border is not the first column of a multi-column character then all the column positions will be replaced with the background character and rendition before the overlay is done. If the default background character is a multi-column character when this occurs, then these subroutines fail.

### Parameters

*srcwin*  
*deswin*

### Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

### Examples

1. To copy *my\_window* on top of *other\_window*, excluding spaces, use:

```
WINDOW *my_window, *other_window;  
overlay(my_window, other_window);
```

2. To copy *my\_window* on top of *other\_window*, including spaces, use:

```
WINDOW *my_window, *other_window;  
overwrite(my_window, other_window);
```

## Related Information

The **copywin** (“copywin Subroutine” on page 543) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Window Data with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## pair\_content Subroutine

### Purpose

Returns the colors in a color pair.

### Library

Curses Library (**libcurses.a**)

### Curses Syntax

```
#include <curses.h>
```

```
pair_content ( Pair, F, B)  
short Pair;  
short *F, *B;
```

### Description

The **pair\_content** subroutine returns the colors in a color pair. A color pair is made up of a foreground and background color. You must call the **start\_color** subroutine before calling the **pair\_content** subroutine.

**Note:** The color pair must already be initialized before calling the **pair\_content** subroutine.

### Return Values

**OK** Indicates the subroutine completed successfully.

**ERR** Indicates the pair has not been initialized.

### Parameters

*Pair* Identifies the color-pair number. The *Pair* parameter must be between **1** and **COLORS\_PAIRS-1**.

*F* Points to the address where the foreground color will be stored. The *F* parameter will be between **0** and **COLORS-1**.

*B* Points to the address where the background color will be stored. The *B* parameter will be between **0** and **COLORS-1**.

### Example

To obtain the foreground and background colors for color-pair 5, use:

```
short *f, *b;  
pair_content(5,f,b);
```

For this subroutine to succeed, you must have already initialized the color pair. The foreground and background colors will be stored at the locations pointed to by *f* and *b*.

## Related Information

The **start\_color** (“start\_color Subroutine” on page 636) subroutine, **init\_pair** (“init\_pair Subroutine” on page 576) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Video Attributes, Working with Color in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## prefresh or pnoutrefresh Subroutine

### Purpose

Updates the terminal and curscr (current screen) to reflect changes made to a pad.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
prefresh(Pad, PY, PX, TTY, TTX, TBY, TBX)
```

```
WINDOW *Pad;
```

```
int PY, PX, TTY;
```

```
int TTX, TBY, TBX;
```

```
pnoutrefresh(Pad, PY, PX, TTY, TTX, TBY, TBX)
```

```
WINDOW *Pad;
```

```
int PY, PX, TTY;
```

```
int TTX, TBY, TBX;
```

### Description

The **prefresh** and **pnoutrefresh** subroutines are similar to the **wrefresh** (“refresh or wrefresh Subroutine” on page 608) and **wnoutrefresh** (“doupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutines. They are different in that pads, instead of windows, are involved, and additional parameters are necessary to indicate what part of the pad and screen are involved.

The *PX* and *PY* parameters specify the upper left corner, in the pad, of the rectangle to be displayed. The *TTX*, *TTY*, *TBX*, and *TBY* parameters specify the edges, on the screen, for the rectangle to be displayed in. The lower right corner of the rectangle to be displayed is calculated from the screen coordinates, since both rectangle and pad must be the same size. Both rectangles must be entirely contained within their respective structures.

The **prefresh** subroutine copies the specified portion of the pad to the physical screen. if you wish to output several pads at once, call **pnoutrefresh** for each pad and then issue one call to **doupdate**. This updates the physical screen once.

### Parameters

*Pad* Specifies the pad to be refreshed.

*PX* (Pad's x-coordinate) Specifies the upper-left column coordinate, in the pad, of the rectangle to be displayed.

*PY* (Pad's y-coordinate) Specifies the upper-left row coordinate, in the pad, of the rectangle to be displayed.

*TBX* (Terminal's Bottom x-coordinate) Specifies the lower-right column coordinate, on the terminal, for the pad to be displayed in.

- TBY** (Terminal's Bottom y-coordinate) Specifies the lower-right row coordinate, on the terminal, for the pad to be displayed in.
- TTX** (Terminal's Top x-coordinate) Specifies the upper-left column coordinate, on the terminal, for the pad to be displayed in.
- TTY** (Terminal's Top Y coordinate) Specifies the upper-left row coordinate, on the terminal, for the pad to be displayed in.

## Examples

1. To update the user-defined `my_pad` pad from the upper-left corner of the pad on the terminal with the upper-left corner at the coordinates `Y=20, X=10` and the lower-right corner at the coordinates `Y=30, X=25` enter

```
WINDOW *my_pad;
prefresh(my_pad, 0, 0, 20, 10, 30, 25);
```

2. To update the user-defined `my_pad1` and `my_pad2` pads and output them both to the terminal in one burst of output, enter:

```
WINDOW *my_pad1; *my_pad2; pnoutrefresh(my_pad1, 0, 0, 20, 10, 30, 25);
pnoutrefresh(my_pad2, 0, 0, 0, 0, 10, 5);
doupdate();
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Window Data with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## printw, wprintw, mvprintw, or mvwprintw Subroutine

### Purpose

Performs a **printf** command on a window using the specified format control string.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
printw( Format, [ Argument ...])
char *Format, *Argument;
```

```
wprintw( Window, Format, [Argument ...])
WINDOW *Window;
char *Format, *Argument;
```

```
mvprintw( Line, Column, Format, [Argument ...])
int Line, Column;
char *Format, *Argument;
```

```
mvprintw(Window, Line, Column, Format, [Argument ...])
```

```
WINDOW *Window;  
int Line, Column;  
char *Format, *Argument;
```

## Description

The **printw**, **wprintw**, **mvprintw**, and **mvprintw** subroutines perform output on a window by using the specified format control string. However, the **waddch** (“addch, mvaddch, mvwaddch, or waddch Subroutine” on page 523) subroutine is used to output characters in a given window instead of invoking the **printf** subroutine. The **mvprintw** and **mvprintw** subroutines move the logical cursor before performing the output.

Use the **printw** and **mvprintw** subroutines on the stdscr and the **wprintw** and **mvprintw** subroutines on user-defined windows.

**Note:** The maximum length of the format control string after expansion is 512 bytes.

## Parameters

<i>Argument</i>	Specifies the item to print. See the <b>printf</b> subroutine for more details.
<i>Column</i>	Specifies the horizontal position to move the cursor to before printing.
<i>Format</i>	Specifies the format for printing the <i>Argument</i> parameter. See the <b>printf</b> subroutine.
<i>Line</i>	Specifies the vertical position to move the cursor to before printing.
<i>Window</i>	Specifies the window to print into.

## Examples

1. To print the user-defined integer variables *x* and *y* as decimal integers in the stdscr, enter:

```
int x, y;  
printw("%d%d", x, y);
```

2. To print the user-defined integer variables *x* and *y* as decimal integers in the user-defined window *my\_window*, enter:

```
int x, y;  
WINDOW *my_window;  
wprintw(my_window, "%d%d", x, y);
```

3. To move the logical cursor to the coordinates *y* = 5, *x* = 10 before printing the user-defined integer variables *x* and *y* as decimal integers in the stdscr, enter:

```
int x, y;  
mvprintw(5, 10, "%d%d", x, y);
```

4. To move the logical cursor to the coordinates *y* = 5, *x* = 10 before printing the user-defined integer variables *x* and *y* as decimal integers in the user-defined window *my\_window*, enter:

```
int x, y;  
WINDOW *my_window;  
mvwprintw(my_window, 5, 10, "%d%d", x, y);
```

## Related Information

The **waddch** (“addch, mvaddch, mvwaddch, or waddch Subroutine” on page 523) subroutine, **printf** subroutine.

The **printf** command.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## putp, tputs Subroutine

### Purpose

Outputs commands to the terminal.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>

int putp(const char *str);

int tputs(const char *str,
int affcnt,
int (*putfunc)(int));
```

### Description

These subroutines output commands contained in the terminfo database to the terminal.

The **putp** subroutine is equivalent to **tputs**(str, 1, putchar). The output of the **putp** subroutine always goes to stdout, not to the files specified in the **setupterm** subroutine.

The **tputs** subroutine outputs *str* to the terminal. The *str* argument must be a terminfo string variable or the return value from the **tgetstr**, **tgoto**, **tigestr**, or **tparm** subroutines. The *affcnt* argument is the number of lines affected, or 1 if not applicable. If the terminfo database indicates that the terminal in use requires padding after any command in the generated string, the **tputs** subroutine inserts pad characters into the string that is sent to the terminal, at positions indicated by the terminfo database. The **tputs** subroutine outputs each character of the generated string by calling the user-supplied **putfunc** subroutine (see below).

The user-supplied **putfunc** subroutine (specified as an argument to the **tputs** subroutine) is either **putchar** or some other subroutine with the same prototype. The **tputs** subroutine ignores the return value of the **putfunc** subroutine.

### Parameters

*\*str*  
*affcnt*  
*\*putfunc*

### Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

### Examples

For the **putp** subroutine:

To call the **tputs**(my\_string, 1, putchar) subroutine, enter:

```
char *my_string;
putp(my_string);
```

For the **tputs** subroutine:

1. To output the clear screen sequence using the user-defined **putchar**-like subroutine **my\_putchar**, enter:

```
int_my_putchar();
tputs(clear_screen, 1 ,my_putchar);
```

2. To output the escape sequence used to move the cursor to the coordinates x=40, y=18 through the user-defined **putchar**-like subroutine **my\_putchar**, enter:

```
int_my_putchar();
tputs(tparm(cursor_address, 18, 40), 1, my_putchar);
```

## Related Information

The **doupdate** (“doupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine, **is\_linetouched** (“is\_linetouched, is\_wintouched, touchline, touchwin, untouchwin, or wtouchin Subroutine” on page 648) subroutine, **putchar** subroutine, **tgetent** (“tgetent, tgetflag, tgetnum, tgetstr, or tgoto Subroutine” on page 640) subroutine, **tigetflag** (“tigetflag, tigetnum, tigetstr, or tparm Subroutine” on page 644) subroutine, **tputs** (“putp, tputs Subroutine” on page 606) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## raw or noraw Subroutine

### Purpose

Places the terminal into or out of raw mode.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
raw( )
noraw( )
```

### Description

The **raw** or **noraw** subroutine places the terminal into or out of raw mode, respectively. RAW mode is similar to CBREAK mode (**cbreak** or **nocbreak** (“cbreak, nocbreak, noraw, or raw Subroutine” on page 535) subroutine). In RAW mode, the system immediately passes typed characters to the user program. The interrupt, quit, and suspend characters are passed uninterrupted, instead of generating a signal. RAW mode also causes 8-bit input and output.

To get character-at-a-time input without echoing, call the **cbreak** and **noecho** subroutines. Most interactive screen-oriented programs require this sort of input.

### Return Values

**OK** Indicates the subroutine completed. The **raw** and **noraw** routines always return this value.

## Examples

1. To place the terminal into raw mode, use:  
`raw();`
2. To place the terminal out of raw mode, use:  
`noraw();`

## Related Information

The **getch** (“getch, mvgetch, mvwgetch, or wgetch Subroutine” on page 561) subroutine, **cbreak** or **nocbreak** (“cbreak, nocbreak, noraw, or raw Subroutine” on page 535) subroutine

Curses Overview for Programming, List of Curses Subroutines, Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## refresh or wrefresh Subroutine

### Purpose

Updates the terminal’s display and the cursor to reflect changes made to a window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
refresh( )
```

```
wrefresh( Window)
WINDOW *Window;
```

### Description

The **refresh** or **wrefresh** subroutines update the terminal and the cursor to reflect changes made to a window. The **refresh** subroutine updates the stdscr. The **wrefresh** subroutine refreshes a user-defined window.

Other subroutines manipulate windows but do not update the terminal’s physical display to reflect their changes. Use the **refresh** or **wrefresh** subroutines to update a terminal’s display after internal window representations change. Both subroutines check for possible scroll errors at display time.

**Note:** The physical terminal cursor remains at the location of the window’s cursor during a refresh, unless the **leaveok** (“leaveok Subroutine” on page 585) subroutine is enabled.

The **refresh** and **wrefresh** subroutines call two other subroutines to perform the refresh operation. First, the **wnoutrefresh** (“doupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine copies the designated window structure to the terminal. Then, the **doupdate** (“doupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine updates the terminal’s display and the cursor.

### Parameters

*Window*      Specifies the window to refresh.

## Examples

1. To update the terminal's display and the current screen structure to reflect changes made to the standard screen structure, use:  

```
refresh();
```
2. To update the terminal and the current screen structure to reflect changes made to a user-defined window called `my_window`, use:  

```
WINDOW *my_window;  
wrefresh(my_window);
```
3. To restore the terminal to its state at the last refresh, use:  

```
wrefresh(curscr);
```

This subroutine is useful if the terminal becomes garbled for any reason.

## Related Information

The **douupdate** (“doupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine, **leaveok** (“leaveok Subroutine” on page 585) subroutine, **wnoutrefresh** (“leaveok Subroutine” on page 585) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## reset\_prog\_mode Subroutine

### Purpose

Restores the terminal to program mode.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>  
reset_prog_mode( )
```

### Description

The **reset\_prog\_mode** subroutine restores the terminal to program or *in curses* mode.

The **reset\_prog\_mode** subroutine is a low-level routine and normally would not be called directly by a program.

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## reset\_shell\_mode Subroutine

### Purpose

Restores the terminal to shell mode.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
reset_shell_mode( )
```

### Description

The **reset\_shell\_mode** subroutine restores the terminal into shell , or "out of curses," mode. This happens automatically when the **endwin** subroutine is called.

### Related Information

The **endwin** ("endwin Subroutine" on page 554) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## resetterm Subroutine

### Purpose

Resets terminal modes to what they were when the **saveterm** subroutine was last called.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
resetterm( )
```

### Description

The **resetterm** subroutine resets terminal modes to what they were when the **saveterm** subroutine was last called.

The **resetterm** subroutine is called by the **endwin** ("endwin Subroutine" on page 554) subroutine, and should normally not be called directly by a program.

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## resetty, savetty Subroutine

### Purpose

Saves/restores the terminal mode.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int resetty(void);
```

```
int savetty(void);
```

### Description

The **resetty** subroutine restores the program mode as of the most recent call to the **savetty** subroutine.

The **savetty** subroutine saves the state that would be put in place by a call to the **reset\_prog\_mode** subroutine.

### Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

### Examples

To restore the terminal to the state it was in at the last call to **savetty**, enter:

```
resetty();
```

## Related Information

The **def\_prog\_mode** (“def\_prog\_mode, def\_shell\_mode, reset\_prog\_mode or reset\_shell\_mode Subroutine” on page 545) subroutine, **endwin** (“endwin Subroutine” on page 554) subroutine, **savetty** (“savetty Subroutine” on page 614) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## restartterm Subroutine

### Purpose

Re-initializes the terminal structures after a restore.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
#include <term.h>
```

```
restartterm ( Term, FileNumber, ErrorCode)
char *Term;
int FileNumber;
int *ErrorCode;
```

### Description

The **restartterm** subroutine is similar to the **setupterm** subroutine except that it is called after restoring memory to a previous state. For example, you would call the **restartterm** subroutine after a call to **scr\_restore** if the terminal type has changed. The **restartterm** subroutine assumes that the windows and the input and output options are the same as when memory was saved, but the terminal type and baud rate may be different.

### Parameters

<i>Term</i>	Specifies the terminal name to obtain the terminal for. If 0 is passed for the parameter, the value of the <b>\$TERM</b> environment variable is used.
<i>FileNumber</i>	Specifies the output file's file descriptor (1 equals standard out).
<i>ErrorCode</i>	Specifies a pointer to an integer to return the error code to. If 0, then the <b>restartterm</b> subroutine exits with an error message instead of returning.

### Example

To restart an **aixterm** after a previous memory save and exit on error with a message, enter:

```
restartterm("aixterm", 1, (int*)0);
```

### Prerequisite Information

Curses Overview for Programming and Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs* .

### Related Information

The **setupterm** ("setupterm Subroutine" on page 624) subroutine.

---

## ripoffline Subroutine

### Purpose

Reserves a line for a dedicated purpose.

## Library

Curses Library (**libcurses.a**)

## Syntax

```
#include
<curses.h>

int
ripoffline(int line,
int (*init)(WINDOW *win,
int columns));
```

## Description

The **ripoffline** subroutine reserves a screen line for use by the application.

Any call to the **ripoffline** subroutine must precede the call to the **initscr** or **newterm** subroutine. If line is positive, one line is removed from the beginning of stdscr; if line is negative, one line is removed from the end. Removal occurs during the subsequent call to the **initscr** or **newterm** subroutine. When the subsequent call is made, the subroutine pointed to by *init* is called with two arguments: a WINDOW pointer to the one-line window that has been allocated and an integer with the number of columns in the window. The initialisation subroutine cannot use the LINES and COLS external variables and cannot call the **wrefresh** or **douupdate** subroutine, but may call the **wnoutrefresh** subroutine.

Up to five lines can be ripped off. Calls to the **ripoffline** subroutine above this limit have no effect, but report success.

## Parameters

*line*  
*\*init*  
*columns*  
*\*win*

## Return Values

The **ripoffline** subroutine returns OK.

## Example

To remove three lines from the top of the screen, enter:

```
#include <curses.h>
ripoffline(1,initfunc);
ripoffline(1,initfunc);
ripoffline(1,initfunc);
initscr();
```

## Related Information

The **douupdate** (“douupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine, **slk\_attoff**, **slk\_init** (“slk\_attoff, slk\_attr\_off, slk\_attron, slk\_attrset, slk\_attr\_set, slk\_clear, slk\_color, slk\_init, slk\_label, slk\_noutrefresh, slk\_refresh, slk\_restore, slk\_set, slk\_touch, slk\_wset, Subroutine” on page 627) subroutine, **initscr** (“initscr and newterm Subroutine” on page 577) subroutine, **newterm** (“newterm Subroutine” on page 594) subroutine.

Curses Overview for Programming and List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## savetty Subroutine

### Purpose

Saves the state of the tty modes.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
savetty( )
```

### Description

The **savetty** subroutine saves the current state of the tty modes in a buffer. It saves the current state in a buffer that the **resetty** subroutine then reads to reset the tty state.

The **savetty** subroutine is called by the **initscr** subroutine and normally should not be called directly by the program.

### Related Information

The **initscr** (“initscr and newterm Subroutine” on page 577) subroutine, **resetty** (“resetty, savetty Subroutine” on page 611) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## scanw, wscanw, mvscanw, or mvwscanw Subroutine

### Purpose

Calls the **wgetstr** subroutine on a window and uses the resulting line as input for a scan.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
scanw( Format, Argument1, Argument2, ...)
char *Format, *Argument1, ...;
```

```
wscanw( Window, Format, Argument1, Argument2, ...)
WINDOW *Window;
char *Format, *Argument1, ...;
```

```
mvscanw( Line, Column, Format, Argument1, Argument2, ...)
int Line, Column;
char *Format, *Argument1, ...;
```

```

mvwscanw(Window, Line, Column, Format, Argument1, Argument2, ...)
WINDOW *Window;
int Line, Column;
char *Format, *Argument1, ...;

```

## Description

The **scanw**, **wscanw**, **mvscanw**, and **mvwscanw** subroutines call the **wgetstr** subroutine on a window and use the resulting line as input for a scan. The **mvscanw** and **mvwscanw** subroutines move the cursor before performing the scan function. Use the **scanw** and **mvscanw** subroutines on the `stdscr` and the **wscanw** and **mvwscanw** subroutines on the user-defined window.

## Parameters

<i>Argument</i>	Specifies the input to read.
<i>Column</i>	Specifies the vertical coordinate to move the logical cursor to before performing the scan.
<i>Format</i>	Specifies the conversion specifications to use to interpret the input. For more information about this parameter, see the discussion of the <i>Format</i> parameter in the <b>scanf</b> (“scanf, fscanf, sscanf, or wsscanf Subroutine” on page 109) subroutine.
<i>Line</i>	Specifies the horizontal coordinate to move the logical cursor to before performing the scan.
<i>Window</i>	Specifies the window to perform the scan in. You only need to specify this parameter with the <b>wscanw</b> and <b>mvwscanw</b> subroutines.

## Example

The following shows how to read input from the keyboard using the **scanw** subroutine.

```

int id;
char deptname[25];

mvprintw(5,0,"Enter your i.d. followed by the department name:\n");
refresh();
scanw("%d %s", &id, deptname);
mvprintw(7,0,"i.d.: %d, Name: %s\n", id, deptname);
refresh();

```

## Related Information

The **wgetstr** (“getnstr, getstr, mvgetnstr, mvgetstr, mvwgetnstr, mvwgetstr, wgetnstr, or wgetstr Subroutine” on page 566) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## scr\_dump, scr\_init, scr\_restore, scr\_set Subroutine

### Purpose

File input/output functions.

### Library

Curses Library (**libcurses.a**)

### Syntax

```

#include <curses.h>

int scr_dump
(const char *filename);

```

```
int scr_init
(const char *filename);

int scr_restore
(const char *filename);

int scr_set
(const char *filename);
```

## Description

The **scr\_dump** subroutine writes the current contents of the virtual screen to the file named by *filename* in an unspecified format.

The **scr\_restore** subroutine sets the virtual screen to the contents of the file named by *filename*, which must have been written using the **scr\_dump** subroutine. The next refresh operation restores the screen to the way it looked in the dump file.

The **scr\_init** subroutine reads the contents of the file named by *filename* and uses them to initialize the Curses data structures to what the terminal currently has on its screen. The next refresh operation bases any updates of this information, unless either of the following conditions is true:

- The terminal has been written to since the virtual screen was dumped to *filename*.
- The terminfo capabilities `rmcup` and `nrrmc` are defined for the current terminal.

The **scr\_set** subroutine is a combination of **scr\_restore** and **scr\_init** subroutines. It tells the program that the information in the file named by *filename* is what is currently on the screen, and also what the program wants on the screen. This can be thought of as a screen inheritance function.

## Parameters

*filename*

## Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

## Examples

For the **scr\_dump** subroutine:

To write the contents of the virtual screen to `/tmp/virtual.dump` file, use:

```
scr_dump("/tmp/virtual.dump");
```

For the **scr\_restore** subroutine:

To restore the contents of the virtual screen from the `/tmp/virtual.dump` file and update the terminal screen, use:

```
scr_restore("/tmp/virtual.dump");
doupdate();
```

## Related Information

The **doupdate** (“`doupdate`, `refresh`, `wnoutrefresh`, or `wrefresh` Subroutines” on page 657) subroutine, **endwin** (“`endwin` Subroutine” on page 554) subroutine, **open** subroutine, **read** (“`read`, `readx`, `readv`, `readvx`, or `pread` Subroutine” on page 16) subroutine, **write** (“`write`, `writex`, `writew`, `writewx` or `pwrite` Subroutines” on page 506) subroutine, **scr\_init** (“`scr_init` Subroutine” on page 617) subroutine, **scr\_restore** (“`scr_restore` Subroutine” on page 618) subroutine.

## scr\_init Subroutine

### Purpose

Initializes the curses data structures from a dump file.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
scr_init( Filename)  
char *Filename;
```

### Description

The **scr\_init** subroutine initializes the curses data structures from a dump file. You create dump files with the **scr\_dump** subroutine. If the file's data is valid, the next screen update is based on the contents of the file rather than clearing the screen and starting from scratch. The data is invalid if the **terminfo** database boolean capability **nrrmc** is TRUE or the contents of the terminal differ from the contents of the dump file.

**Note:** If **nrrmc** is TRUE, avoid calling the **putp** subroutine with the **exit\_ca\_mode** value before calling **scr\_init** subroutine in your application.

You can call the **scr\_init** subroutine after the **initscr** subroutine to update the screen with the dump file contents. Using the **keypad**, **meta**, **slk\_clear**, **curs\_set**, **flash**, and **beep** subroutines do not affect the contents of the screen, but cause the terminal's modification time to change.

You can allow more than one process to share screen dumps. Both processes must be run from the same terminal. The **scr\_init** subroutine first ensures that the process that created the dump is in sync with the current terminal data. If the modification time of the terminal is not the same as that specified in the dump file, the **scr\_init** subroutine assumes that the screen image on the terminal has changed from that in the file, and the file's data is invalid.

If you are allowing two processes to share a screen dump, it is important to understand that one process starts up another process. The following activities happen:

- The second process creates the dump file with the **scr\_init** subroutine.
- The second process exits without causing the terminal's time stamp to change by calling the **endwin** subroutine followed by the **scr\_dump** subroutine, and then the **exit** subroutine.
- Control is passed back to the first process.
- The first process calls the **scr\_init** subroutine to update the screen contents with the dump file data.

### Return Values

**ERR** Indicates the dump file's time stamp is old or the boolean capability **nrrmc** is TRUE.

**OK** Indicates that the curses data structures were successfully initialized using the contents of the dump file.

## Parameters

*Filename*        Points to a dump file.

## Related Information

The **scr\_dump** (“scr\_dump, scr\_init, scr\_restore, scr\_set Subroutine” on page 615) subroutine, **scr\_restore** (“scr\_restore Subroutine”) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Window Data with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## scr\_restore Subroutine

### Purpose

Restores the virtual screen from a dump file.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
scr_restore( FileName)  
char *FileName;
```

### Description

The **scr\_restore** subroutine restores the virtual screen from the contents of a dump file. You create a dump file with the **scr\_dump** subroutine. To update the terminal’s display with the restored virtual screen, call the **wrefresh** or **douupdate** subroutine after restoring from a dump file.

To communicate the screen image across processes, use the **scr\_restore** subroutine along with the **scr\_dump** subroutine.

### Return Values

**ERR**    Indicates the content of the dump file is incompatible with the current release of curses.

**OK**     Indicates that the virtual screen was successfully restored from a dump file.

## Parameters

*FileName*        Identifies the name of the dump file.

### Example

To restore the contents of the virtual screen from the `/tmp/virtual.dump` file and update the terminal screen, use:

```
scr_restore("/tmp/virtual.dump");  
douupdate();
```

## Related Information

The **scr\_dump** (“scr\_dump, scr\_init, scr\_restore, scr\_set Subroutine” on page 615) subroutine, **scr\_init** (“scr\_init Subroutine” on page 617) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Understanding Terminals with Curses, Manipulating Video Attributes in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## scr1, scroll, wscr1 Subroutine

### Purpose

Scrolls a Curses window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
int scr1
(int n);
int scroll
(WINDOW *win);

int wscr1
(WINDOW *win,
int n);
```

### Description

The **scroll** subroutine scrolls win one line in the direction of the first line

The **scr1** and **wscr1** subroutines scroll the current or specified window. If *n* is positive, the window scrolls *n* lines toward the first line. Otherwise, the window scrolls *-n* lines toward the last line.

These subroutines do not change the cursor position. If scrolling is disabled for the current or specified window, these subroutines have no effect. The interaction of these subroutines with the **setscreg** subroutine is currently unspecified.

### Parameters

*\*win*     Specifies the window to scroll.  
*n*

### Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

### Examples

To scroll the user-defined window `my_window` up one line, enter:

```
WINDOW *my_window;
scroll(my_window);
```

## Related Information

The **scrollok** (“scrollok Subroutine”) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## scrollok Subroutine

### Purpose

Enables or disables scrolling.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
scrollok( Window, Flag)  
WINDOW *Window;  
bool Flag;
```

### Description

The **scrollok** subroutine enables or disables scrolling. Scrolling occurs when a program or user:

- Moves the cursor off the window’s bottom edge.
- Enters a new-line character on the last line.
- Types the last character of the last line.

If enabled, **curses** calls a refresh as part of the scrolling action on both the window and the physical display. To get the physical scrolling effect on the terminal, it is also necessary to call the **idlok** (“idlok Subroutine” on page 573) subroutine.

If scrolling is disabled, the cursor is left on the bottom line at the location where the character was entered.

### Parameters

*Flag* Enables scrolling when set to TRUE. Otherwise, set the *Flag* parameter to FALSE to disable scrolling.  
*Window* Identifies the window to enable or disable scrolling in.

### Examples

1. To turn scrolling on in the user-defined window *my\_window*, enter:

```
WINDOW *my_window;  
scrollok(my_window, TRUE);
```

2. To turn scrolling off in the user-defined window *my\_window*, enter:

```
WINDOW *my_window;  
scrollok(my_window, FALSE);
```

## Related Information

The **idlok** (“idlok Subroutine” on page 573) subroutine.

## set\_curterm Subroutine

### Purpose

Sets the current terminal variable to the specified terminal.

### Library

Curses Library (**libcurses.a**)

### Curses Syntax

```
#include <curses.h>
#include <term.h>
```

```
set_curterm( Newterm)
TERMINAL *Newterm;
```

### Description

The **cur\_term** subroutine sets the **cur\_term** variable to the terminal specified by the *Newterm* parameter. The **cur\_term** subroutine is useful when the **setupterm** subroutine is called more than once. The **set\_curterm** subroutine allows the programmer to toggle back and forth between terminals.

When information for a particular terminal is no longer required, remove it using the **del\_curterm** subroutine.

**Note:** The **cur\_term** subroutine is a low-level subroutine. You should use this subroutine only if your application must deal directly with the **terminfo** database to handle certain terminal capabilities. For example, use this subroutine if your application programs function keys.

### Parameters

*Newterm*      Points to a **TERMINAL** structure. This structure contains information about a specific terminal.

### Examples

To set the **cur\_term** variable to point to the `my_term` terminal, use:

```
TERMINAL *newterm;
set_curterm(newterm);
```

### Related Information

The **setupterm** (“setupterm Subroutine” on page 624) subroutine.

Curses Overview for Programming and List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## setscrreg or wsetscrreg Subroutine

### Purpose

Creates a software scrolling region within a window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
setscrreg( Tmargin, Bmargin)
int Tmargin, Bmargin;
```

```
wsetscrreg( Window, Tmargin, Bmargin)
WINDOW *Window;
int Tmargin, Bmargin;
```

### Description

The **setscrreg** and **wsetscrreg** subroutines create a software scrolling region within a window. Use the **setscrreg** subroutine with the stdscr and the **wsetscrreg** subroutine with user-defined windows.

You pass the **setscrreg** subroutines values for the top line and bottom line of the region. If the **setscrreg** subroutine and **scrollok** subroutine are enabled for the region, any attempt to move off the line specified by the *Bmargin* parameter causes all the lines in the region to scroll up one line.

**Note:** Unlike the **idlok** subroutine, the **setscrreg** subroutines have nothing to do with the use of a physical scrolling region capability that the terminal may or may not have.

### Parameters

<i>Bmargin</i>	Specifies the last line number in the scrolling region.
<i>Tmargin</i>	Specifies the first line number in the scrolling region (0 is the top line of the window.)
<i>Window</i>	Specifies the window to place the scrolling region in. You specify this parameter only with the <b>wsetscrreg</b> subroutine.

### Examples

1. To set a scrolling region starting at the 10th line and ending at the 30th line in the stdscr, enter:  

```
setscrreg(9, 29);
```

**Note:** Zero is always the first line.

2. To set a scrolling region starting at the 10th line and ending at the 30th line in the user-defined window *my\_window*, enter:  

```
WINDOW *my_window;
wsetscrreg(my_window, 9, 29);
```

### Related Information

The **idlok** (“idlok Subroutine” on page 573) subroutine, **scrollok** (“scrollok Subroutine” on page 620) subroutine, **wrefresh** (“refresh or wrefresh Subroutine” on page 608) subroutine.

## setsyx Subroutine

### Purpose

Sets the coordinates of the virtual screen cursor.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
setsyx( Y, X)  
int Y, X;
```

### Description

The **setsyx** subroutine sets the coordinates of the virtual screen cursor to the specified row and column coordinates. If *Y* and *X* are both -1, then the **leaveok** flag is set. (**leaveok** may be set by applications that do not use the cursor.)

The **setsyx** subroutine is intended for use in combination with the **getsyx** subroutine. These subroutines should be used by a user-defined function that manipulates curses windows but wants the position of the cursor to remain the same. Such a function would do the following:

- Call the **getsyx** subroutine to obtain the current virtual cursor coordinates.
- Continue processing the windows.
- Call the **wnoutrefresh** subroutine on each window manipulated.
- Call the **setsyx** subroutine to reset the current virtual cursor coordinates to the original values.
- Refresh the display by calling the **doupdate** subroutine.

### Parameters

- X* Specifies the column to set the virtual screen cursor to.  
*Y* Specifies the row to set the virtual screen cursor to.

### Related Information

The **doupdate** (“doupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine, **getsyx** (“getsyx Subroutine” on page 568) subroutine, **leaveok** (“leaveok Subroutine” on page 585) subroutine, **wnoutrefresh** (“doupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657) subroutine.

Controlling the Cursor with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Curses Overview for Programming and List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## set\_term Subroutine

### Purpose

Switches between screens.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
SCREEN *set_term  
(SCREEN *new);
```

### Description

The **set\_term** subroutine switches between different screens. The *new* argument specifies the current screen.

### Parameters

*\*new*

### Return Values

Upon successful completion, the **set\_term** subroutine returns a pointer to the previous screen. Otherwise, it returns a null pointer.

### Examples

To make the terminal stored in the user-defined **SCREEN** variable `my_terminal` the current terminal and then store a pointer to the old terminal in the user-defined variable `old_terminal`, enter:

```
SCREEN *old_terminal, *my_terminal;  
old_terminal = set_term(my_terminal);
```

### Related Information

The **initscr** (“initscr and newterm Subroutine” on page 577) subroutine, **newterm** (“newterm Subroutine” on page 594) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## setupterm Subroutine

### Purpose

Initializes the terminal structure with the values in the **terminfo** database.

## Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>
#include <term.h>
```

```
setupterm( Term, FileNumber, ErrorCode)
char *Term;
int FileNumber;
int *ErrorCode;
```

## Description

The **setupterm** subroutine determines the number of lines and columns available on the output terminal. The **setupterm** subroutine calls the **termdef** subroutine to define the number of lines and columns on the display. If the **termdef** subroutine cannot supply this information, the **setupterm** subroutine uses the values in the **terminfo** database.

The **setupterm** subroutine initializes the terminal structure with the terminal-dependent capabilities from **terminfo**. This routine is automatically called by the **initscr** and **newterm** subroutines. The **setupterm** subroutine deals directly with the **terminfo** database.

Two of the terminal-dependent capabilities are the lines and columns. The **setupterm** subroutine populates the lines and column fields in the terminal structure in the following manner:

1. If the environment variables **LINES** and **COLUMNS** are set, the **setupterm** subroutine uses these values.
2. If the environment variables are not set, the **setupterm** subroutine obtains the lines and columns information from the tty subsystem.
3. As a last resort, the **setupterm** subroutine uses the values defined in the **terminfo** database.

**Note:** These may or may not be the same as the values in the **terminfo** database.

The simplest call is **setupterm((char\*) 0, 1, (int\*) 0)**, which uses all defaults.

After the call to the **setupterm** subroutine, the **cur\_term** global variable is set to point to the current structure of terminal capabilities. A program can use more than one terminal at a time by calling the **setupterm** subroutine for each terminal and then saving and restoring the **cur\_term** variable.

## Parameters

<i>ErrorCode</i>	Specifies a pointer to an integer to return the error code to. If a null pointer (0) is passed for this parameter, no status is returned. An error causes the <b>setupterm</b> subroutine to print an error message and exit instead of returning.
<i>FileNumber</i>	Specifies the output files file descriptor (1 equals standard output).
<i>Term</i>	Specifies the terminal name. If 0 is passed for this parameter, the value of the <b>\$TERM</b> environment variable is used.

## Return Values

One of the following status values is stored into the integer pointed to by the *ErrorCode* parameter:

- 1 Successful completion.
- 0 No such terminal.

-1 An error occurred while locating the **terminfo** database.

## Example

To determine the current terminal's capabilities using **\$TERM** as the terminal name, standard output as output, and returning no error codes, enter:

```
setupterm((char*) 0, 1, (int*) 0);
```

## Related Information

The **termdef** ("termdef Subroutine" on page 352) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## **\_showstring** Subroutine

### Purpose

Dumps the string in the specified string address to the terminal at the specified location.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
_showstring(Line, Column, First, Last, String)  
int Line, Column, First, Last;  
char * String;
```

### Description

The **\_showstring** subroutine dumps the string in the specified string address to the terminal at the specified location. This is an internal extended curses subroutine and should not normally be called directly by the program.

### Parameters

<i>Column</i>	Specifies the horizontal coordinate of the terminal at which to dump the string.
<i>First</i>	Specifies the beginning string address of the string to dump to the terminal.
<i>Last</i>	Specifies the end string address of the string to dump to the terminal.
<i>Line</i>	Specifies the vertical coordinate of the terminal at which to dump the string.
<i>String</i>	Specifies the string to dump to the terminal.

## Related Information

Curses Overview for Programming , List of Curses Subroutines , Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

**slk\_attroff, slk\_attr\_off, slk\_attron, slk\_attrset, slk\_attr\_set, slk\_clear, slk\_color, slk\_init, slk\_label, slk\_noutrefresh, slk\_refresh, slk\_restore, slk\_set, slk\_touch, slk\_wset, Subroutine**

## Purpose

Soft label subroutines.

## Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>
```

```
int slk_attroff  
(const chtype attrs);
```

```
int slk_attr_off  
(const attr_t attrs,  
void *opts);
```

```
int slk_attron  
(const chtype attrs);
```

```
int slk_attr_on  
(const attr_t attrs,  
void *opts);
```

```
int slk_attrset  
(const chtype attrs);
```

```
int slk_attr_set  
(const attr_t attrs,  
short color_pair_number,  
void *opts);
```

```
int slk_clear  
(void);
```

```
int slk_color  
(short color_pair_number);
```

```
int slk_init  
(int fmt);
```

```
char *slk_label  
(int labnum);
```

```
int slk_noutrefresh  
(void);
```

```
int slk_refresh  
(void);
```

```
int slk_restore  
(void);
```

```
int slk_set  
(int labnum,  
const char *label,  
int justify);
```

```
int slk_touch
(void);

int slk_wset
(int labnum,
const wchar_t *label,
int justify);
```

## Description

The Curses interface manipulates the set of soft function-key labels that exist on many terminals. For those terminals that do not have sort labels, Curses takes over the bottom line of *stdscr*, reducing the size of *stdscr* and the value of the LINES external variable. There can be up to eight labels of up to eight display columns each.

To use soft labels, the **slk\_init** subroutine must be called before **initscr**, **newterm**, or **ripoffline** is called. If **initscr** eventually uses a line from *stdscr* to emulate the soft labels, then *fmt* determines how the labels are arranged on the screen. Setting *fmt* to **0** indicates a 3-2-3 arrangement of the labels; **1** indicates a 4-4 arrangement. Other values for *fmt* are unspecified.

The **slk\_init** subroutine has the effect of calling the **ripoffline** subroutine to reserve one screen line to accommodate the requested format.

The **slk\_set** and **slk\_wset** subroutines specify the text of soft label number *labnum*, within the range from 1 to and including 8. The *label* argument is the string to be put on the label. With **slk\_set** and **slk\_wset**, the width of the label is limited to eight column positions. A null string or a null pointer specifies a blank label. The *justify* argument can have the following values to indicate how to justify label within the space reserved for it:

- 0**     Align the start of label with the start of the space.
- 1**     Center label within the space.
- 2**     Align the end of label with the end of the space.

The **slk\_refresh** and **slk\_noutrefresh** subroutines correspond to the **wrefresh** and **wnoutrefresh** subroutines.

The **slk\_label** subroutine obtains soft label number *labnum*.

The **slk\_clear** subroutine immediately clears the soft labels from the screen.

The **slk\_touch** subroutine forces all the soft labels to be output the next time **slk\_noutrefresh** or **slk\_refresh** subroutines is called.

The **slk\_attron**, **slk\_attrset** and **slk\_attroff** subroutines correspond to the **attron**, **attrset**, and **attroff** subroutines. They have an effect only if soft labels are simulated on the bottom line of the screen.

The **slk\_attr\_off**, **slk\_attr\_on**, **slk\_sttr\_set**, and **slk\_attroff** subroutines correspond to the **slk\_attron**, **slk\_attrset**, and **color\_set** and thus support the attribute constants with the WA\_prefix and color.

The *opts* argument is reserved for definition in a future edition of this document. Currently, the application must provide a null pointer as *opts*.

## Parameters

*attrs*  
*\*opts*

*color\_pair\_number*  
*fmt*  
*labnum*  
*justify*  
*\*label*

## Examples

For the **slk\_init** subroutine:

To initialize soft labels on a terminal that does not support soft labels internally, do the following:

```
slk_init(1);
```

This example arranges the labels so that four labels appear on the right of the screen and four appear on the left.

For the **slk\_label** subroutine:

To obtain the label name for soft label 3, use:

```
char *label_name;  
label_name = slk_label(3);
```

For the **slk\_noutrefresh** subroutine:

To refresh soft label 8 on the virtual screen but not on the physical screen, use:

```
slk_set(8, "Insert", 1);  
slk_noutrefresh();
```

For the **slk\_refresh** subroutine:

To set and left-justify the soft labels and then refresh the physical screen, use:

```
slk_init(0);  
initscr();  
slk_set(1, "Insert", 0);  
slk_set(2, "Quit", 0);  
slk_set(3, "Add", 0);  
slk_set(4, "Delete", 0);  
slk_set(5, "Undo", 0);  
slk_set(6, "Search", 0);  
slk_set(7, "Replace", 0);  
slk_set(8, "Save", 0);  
slk_refresh();
```

For the **slk\_set** subroutine:

```
slk_set(2, "Quit", 1);
```

## Return Values

Upon successful completion, the **slk\_label** subroutine returns the requested label with leading and trailing blanks stripped. Otherwise, it returns a null pointer.

Upon successful completion, the other subroutines return OK. Otherwise, they return ERR.

## Related Information

The **attroff** (“attroff, attron, attrset, wattroff, wattron, or wattrset Subroutine” on page 526) subroutine, **ripoffline** (“ripoffline Subroutine” on page 612) subroutine, **wcswidth** (“wcswidth Subroutine” on page 465) subroutine, **slk\_init** (“slk\_init Subroutine”) subroutine, **slk\_set** (“slk\_set Subroutine” on page 633) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Video Attributes in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## slk\_init Subroutine

### Purpose

Initializes soft function-key labels.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
slk_init( Labfmt)  
int Labfmt;
```

### Description

The **slk\_init** subroutine initializes soft function-key labels. This is one of several subroutines curses provides for manipulating soft function-key labels. These labels appear at the bottom of the screen and give applications, such as editors, a more user-friendly look. To use soft labels, you must call the **slk\_init** subroutine before calling the **initscr** or **newterm** subroutine.

Some terminals support soft labels, others do not. For terminals that do not support soft labels, Curses emulates soft labels by using the bottom line of the stdscr. To accommodate soft labels, curses reduces the size of the stdscr and the **LINES** environment variable as required.

### Parameter

*Labfmt* Simulates soft labels. To arrange three labels on the right, two in the center, and three on the right of the screen, specify a 0 for this parameter. To arrange four labels on the left and four on the right of the screen, specify a 1 for this parameter.

### Example

To initialize soft labels on a terminal that does not support soft labels internally, do the following:

```
slk_init(1);
```

This example arranges the labels so that four labels appear on the right of the screen and four appear on the left.

## Related Information

The **initscr** (“initscr and newterm Subroutine” on page 577) subroutine, **newterm** (“newterm Subroutine” on page 594) subroutine.

## slk\_label Subroutine

### Purpose

Returns the label name for a specified soft label.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
char *slk_label( LabNum)
int LabNum;
```

### Description

The **slk\_label** subroutine returns the label name for a specified soft function-key label. These labels appear at the bottom of the screen and give applications, such as editors, a more user-friendly look. The **slk\_label** subroutine returns the name in the format it was in when passed to the **slk\_set** subroutine. If the name was justified by the **slk\_set** subroutine, the justification is removed.

### Parameters

*LabNum* Specifies the label number. This parameter must be in the range 1 to 8.

### Example

To obtain the label name for soft label 3, use:

```
char *label_name;
label_name = slk_label(3);
```

### Return Values

**NULL** Indicates a label number that is not valid or a label number not set with the **slk\_set** subroutine.

**OK** Indicates that the label name was successfully retrieved.

### Related Information

The **slk\_init** (“slk\_init Subroutine” on page 630) subroutine and **slk\_set** (“slk\_set Subroutine” on page 633) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Video Attributes in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## slk\_noutrefresh Subroutine

### Purpose

Updates the soft labels on the virtual screen.

## Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>
slk_noutrefresh()
```

## Description

The **slk\_noutrefresh** subroutine updates the soft function-key labels on the virtual screen. These labels appear at the bottom of the screen and give applications, such as editors, a more user-friendly look. This subroutine is useful for updating multiple labels. You can use the **slk\_noutrefresh** subroutine to update all soft labels on the virtual screen with no updates to the physical screen. To update the physical screen, use the **slk\_refresh** or **refresh** subroutine.

## Example

To refresh soft label 8 on the virtual screen but not on the physical screen, use:

```
slk_set(8, "Insert", 1);
slk_noutrefresh();
```

## Related Information

The **slk\_init** (“slk\_init Subroutine” on page 630) subroutine, **slk\_refresh** (“slk\_refresh Subroutine”) subroutine, **wrefresh** (“refresh or wrefresh Subroutine” on page 608) subroutine.

Curses Overview for Programming, Manipulating Video Attributes, List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## slk\_refresh Subroutine

### Purpose

Updates soft labels on the virtual and physical screens.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
slk_refresh()
```

### Description

The **slk\_refresh** subroutine refreshes the virtual and physical screens after an update to soft function-key labels. These labels appear at the bottom of the screen and give applications, such as editors, a more user-friendly look.

### Example

To set and left-justify the soft labels and then refresh the physical screen, use:

```
slk_init(0);
initscr();
slk_set(1, "Insert", 0);
slk_set(2, "Quit", 0);
slk_set(3, "Add", 0);
slk_set(4, "Delete", 0);
```

```
slk_set(5, "Undo", 0);
slk_set(6, "Search", 0);
slk_set(7, "Replace", 0);
slk_set(8, "Save", 0);
slk_refresh();
```

## Related Information

The **slk\_init routine** (“slk\_init Subroutine” on page 630) subroutine, **slk\_set routine** (“slk\_set Subroutine”) subroutine, **slk\_noutrefresh** (“slk\_noutrefresh Subroutine” on page 631) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Video Attributes in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## slk\_restore Subroutine

### Purpose

Restores soft function-key labels to the screen.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>

slk_restore()
```

### Description

The **slk\_restore** subroutine restores the soft function-key labels to the screen after a call to the **slk\_clear** subroutine. The label names are not restored. These labels appear at the bottom of the screen and give applications, such as editors, a more user-friendly look. You must call the **slk\_init** subroutine before you can use soft labels.

## Related Information

The **slk\_init** (“slk\_init Subroutine” on page 630) subroutine, **slk\_clear** (“slk\_attron, slk\_attr\_off, slk\_attron, slk\_attrset, slk\_attr\_set, slk\_clear, slk\_color, slk\_init, slk\_label, slk\_noutrefresh, slk\_refresh, slk\_restore, slk\_set, slk\_touch, slk\_wset, Subroutine” on page 627) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Soft Labels in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## slk\_set Subroutine

### Purpose

Sets up soft function-key labels.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
slk_set(LabNum, LabStr, LabFmt)
int LabNum;
char * LabStr;
int LabFmt;
```

## Description

The **slk\_set** subroutine sets up each soft function-key label with the appropriate name. These labels appear at the bottom of the screen and give applications, such as editors, a more user-friendly look. Label names are restricted to 8 characters each.

## Parameters

<i>LabNum</i>	Specifies the label number. The value can range from 1 to 8.
<i>LabStr</i>	Specifies the string (name) to put on the label. If the string is NULL, the label is blank.
<i>LabFmt</i>	Specifies the label alignment. The following values are valid:
0	Left-justified
1	Centered
2	Right-justified

## Example

```
slk_set(2, "Quit", 1);
```

## Related Information

The **slk\_init** (“slk\_init Subroutine” on page 630) routine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Video Attributes in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## slk\_touch Subroutine

### Purpose

Forces an update of the soft function-key labels.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
slk_touch()
```

### Description

The **slk\_touch** subroutine forces an update of the soft function-key labels on the physical screen the next time the **slk\_noutrefresh** subroutine is called. These labels appear at the bottom of the screen and give applications, such as editors, a more user-friendly look. You must call the **slk\_init** subroutine before using soft labels.

## Related Information

The **slk\_init** (“slk\_init Subroutine” on page 630) subroutine.

## standend, standout, wstandend, or wstandout Subroutine

### Purpose

Sets and clears window attributes.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int standend  
(void);
```

```
int standout  
(void);
```

```
int wstandend  
(WINDOW *win);
```

```
int wstandout  
(WINDOW *win);
```

### Description

The **standend** and **standout** subroutines turn off all attributes of the current or specified window.

The **wstandout** and **wstandend** subroutines turn on the **standout** attribute of the current or specified window.

### Parameters

*\*win* Specifies the window in which to set the attributes.

### Return Values

These subroutines always return 1.

### Examples

1. To turn on the **standout** attribute in the stdscr, enter:

```
standout();
```

This example is functionally equivalent to:

```
attron(A_STANDOUT);
```

2. To turn on the **standout** attribute in the user-defined window `my_window`, enter:

```
WINDOW *my_window;  
wstandout(my_window);
```

This example is functionally equivalent to:

```
wattron(my_window, A_STANDOUT);
```

3. To turn off the **standout** attribute in the default window, enter:

```
standend();
```

This example is functionally equivalent to:

```
attroff(A_STANDOUT);
```

4. To turn off the **standout** attribute in the user-defined window `my_window`, enter:

```
WINDOW *my_window;  
wstandend(my_window);
```

This example is functionally equivalent to:

```
wattroff(my_window, A_STANDOUT);
```

## Related Information

The **attroff**, **attron**, or **wattroff** (“attroff, attron, attrset, wattroff, wattron, or wattrset Subroutine” on page 526) subroutines.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Video Attributes in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## start\_color Subroutine

### Purpose

Initializes color.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>  
start_color();
```

### Description

The **start\_color** subroutine initializes color. This subroutine requires no arguments. You must call the **start\_color** subroutine if you intend to use color in your application. Except for the **has\_colors** and **can\_change\_color** subroutines, you must call the **start\_color** subroutine before any other color manipulation subroutine. A good time to call **start\_color** is right after calling the **initscr** routine and after establishing whether the terminal supports color.

The **start\_color** routine initializes the following basic colors:

<b>COLOR_BLACK</b>	0
<b>COLOR_BLUE</b>	1
<b>COLOR_GREEN</b>	2
<b>COLOR_CYAN</b>	3
<b>COLOR_RED</b>	4
<b>COLOR_MAGENTA</b>	5
<b>COLOR_YELLOW</b>	6
<b>COLOR_WHITE</b>	7

The subroutine also initializes two global variables: **COLORS** and **COLOR\_PAIRS**. The **COLORS** variable is the maximum number of colors supported by the terminal. The **COLOR\_PAIRS** variable is the maximum number of color-pairs supported by the terminal.

The **start\_color** subroutine also restores the terminal's colors to the original values right after the terminal was turned on.

## Return Values

**ERR** Indicates the terminal does not support colors.

**OK** Indicates the terminal does support colors.

## Example

To enable the color support for a terminal that supports color, use:

```
start_color();
```

## Related Information

The **has\_colors** (“has\_colors Subroutine” on page 570) subroutine, **can\_change\_color** (“can\_change\_color, color\_content, has\_colors,init\_color, init\_pair, start\_color or pair\_content Subroutine” on page 532) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Manipulating Video Attributes in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## subpad Subroutine

### Purpose

Creates a subwindow within a pad.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
WINDOW *subpad(Orig, NLines, NCols, Begin_Y, Begin_X)
```

```
WINDOW * Orig;
```

```
int NCols, NLines, Begin_Y, Begin_X;
```

### Description

The **subpad** subroutine creates and returns a pointer to a subpad. A subpad is a window within a pad. You specify the size of the subpad by supplying a starting coordinate and the number of rows and columns within the subpad. Unlike the **subwin** subroutine, the starting coordinates are relative to the pad and not the terminal's display.

Changes to the subpad affect the character image of the parent pad, as well. If you change a subpad, use the **touchwin** or **touchline** subroutine on the parent pad before refreshing the parent pad. Use the **prefresh** subroutine to refresh a pad.

## Parameters

<i>Orig</i>	Points to the parent pad.
<i>NLines</i>	Specifies the number of lines (rows) in the subpad.
<i>NCols</i>	Specifies the number of columns in the subpad.
<i>Begin_Y</i>	Identifies the upper left-hand row coordinate of the subpad relative to the parent pad.
<i>Begin_X</i>	Identifies the upper left-hand column coordinate of the subpad relative to the parent pad.

## Examples

To create a subpad, use:

```
WINDOW *orig, *mypad;
orig = newpad(100, 200);
mypad = subpad(orig, 30, 5, 25, 180);
```

The parent pad is 100 lines by 200 columns. The subpad is 30 lines by 5 columns and starts in line 25, column 180 of the parent pad.

## Related Information

Curses Overview for Programming, List of Curses Subroutines, Windows in the Curses Environment in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## subwin Subroutine

### Purpose

Creates a subwindow within an existing window.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
WINDOW *subwin (ParentWindow, NumLines, NumCols, Line, Column)
WINDOW * ParentWindow ;
int NumLines, NumCols, Line, Column;
```

### Description

The **subwin** subroutine creates a subwindow within an existing window. You must supply coordinates for the subwindow relative to the terminal's display. Recall that the subwindow shares its parent's window buffer. Changes made to the shared window buffer in the area covered by a subwindow, through either the parent window or any of its subwindows, affects all windows sharing the window buffer.

When changing the image of a subwindow, it is necessary to call the **touchwin** ("touchwin Subroutine" on page 650) or **touchline** subroutine on the parent window before calling the **wrefresh** ("refresh or wrefresh Subroutine" on page 608) subroutine on the parent window.

Changes to one window will affect the character image of both windows.

## Parameters

<i>NumCols</i>	Indicates the number of vertical columns in the subwindow's width. If 0 is passed as the <i>NumCols</i> value, the subwindow runs from the Column to the right edge of its parent window.
----------------	---

<i>NumLines</i>	Indicates the number of horizontal lines in the subwindow's height. If 0 is passed as the <i>NumLines</i> parameter, then the subwindow runs from the Line to the bottom of its parent window.
<i>ParentWindow</i>	Specifies the subwindow's parent.
<i>Column</i>	Specifies the horizontal coordinate for the upper-left corner of the subwindow. This coordinate is relative to the (0, 0) coordinates of the terminal, not the (0, 0) coordinates of the parent window. <b>Note:</b> The upper-left corner of the terminal is referenced by the coordinates (0, 0).
<i>Line</i>	Specifies the vertical coordinate for the upper-left corner of the subwindow. This coordinate is relative to the (0, 0) coordinates of the terminal, not the (0, 0) coordinates of the parent window. <b>Note:</b> The upper-left corner of the terminal is referenced by the coordinates (0, 0).

## Return Values

When the **subwin** subroutine is successful, it returns a pointer to the subwindow structure. Otherwise, it returns the following:

**ERR** Indicates one or more of the parameters is invalid or there is insufficient storage available for the new structure.

## Examples

1. To create a subwindow, use:

```
WINDOW *my_window, *my_sub_window;
```

```
my_window = newwin ("derwin, newwin, or subwin Subroutine" on page 596)
              (5, 10, 20, 30);
```

```
my_sub_window = subwin(my_window, 2, 5, 20, 30);
```

*my\_sub\_window* is now a subwindow 2 lines deep, 5 columns wide, starting at the same coordinates of its parent window *my\_window*. That is, the subwindow's upper-left corner is at coordinates  $y = 20$ ,  $x = 30$  and lower-right corner is at coordinates  $y = 21$ ,  $x = 34$ .

2. To create a subwindow that is flush with the right side of its parent, use:

```
WINDOW *my_window, *my_sub_window;
```

```
my_window = newwin ("derwin, newwin, or subwin Subroutine" on page 596)
              (5, 10, 20, 30);
```

```
my_sub_window = subwin(my_window, 2, 0, 20, 30);
```

*my\_sub\_window* is now a subwindow 2 lines deep, extending all the way to the right side of its parent window *my\_window*, and starting at the same coordinates. That is, the subwindow's upper-left corner is at coordinates  $y = 20$ ,  $x = 30$  and lower-right corner is at coordinates  $y = 21$ ,  $x = 39$ .

3. To create a subwindow in the lower-right corner of its parent, use:

```
WINDOW *my_window, *my_sub_window
```

```
my_window = newwin ("derwin, newwin, or subwin Subroutine" on page 596)
              (5, 10, 20, 30);
```

```
my_sub_window = subwin(my_window, 0, 0, 22, 35);
```

`my_sub_window` is now a subwindow that fills the bottom right corner of its parent window, `my_window`, starting at the coordinates `y = 22, x = 35`. That is, the subwindow's upper-left corner is at coordinates `y = 22, x = 35` and lower-right corner is at coordinates `y = 24, x = 39`.

## Related Information

The **touchwin** (“touchwin Subroutine” on page 650), **newwin** (“derwin, newwin, or subwin Subroutine” on page 596), and **wrefresh** (“refresh or wrefresh Subroutine” on page 608) subroutines.

Curses Overview for Programming, List of Curses Subroutines, Windows in the Curses Environment in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tgetent, tgetflag, tgetnum, tgetstr, or tgoto Subroutine

### Purpose

Termcap database emulation.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int tgetent  
(char *bp,  
const char *name);
```

```
int tgetflag  
(char id[2]);
```

```
int tgetnum  
(char id[2]);
```

```
char *tgetstr  
(char id[2],  
char **area);
```

```
char *tgoto  
(char *cap,  
int col,  
int row);
```

### Description

The **tgetent** subroutine looks up the termcap entry for *name*. The emulation ignores the buffer pointer *bp*.

The **tgetflag** subroutine gets the boolean entry for *id*.

The **tgetnum** subroutine gets the numeric entry for *id*.

The **tgetstr** subroutine gets the string entry for *id*. If *area* is not a null pointer and does not point to a null pointer, the **tgetstr** subroutine copies the string entry into the buffer pointed to by *\*area* and advances the variable pointed to by *area* to the first byte after the copy of the string entry.

The **tgoto** subroutine instantiates the parameters *col* and *row* into the capability *cap* and returns a pointer to the resulting string.

All of the information available in the terminfo database need not be available through these subroutines.

## Parameters

*bp*  
*name*  
*col*  
*row*  
*\*\*area*  
*cap*        *id[2]*

## Return Values

Upon successful completion, subroutines that return an integer return OK. Otherwise, they return ERR.

## Related Information

The **putc**, **setupterm** (“setupterm Subroutine” on page 624), **tigetflg** (“tigetflag, tigetnum, tigetstr, or tparm Subroutine” on page 644) subroutines.

Curses Overview for Programming, List of Curses Subroutines, Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tigetflag Subroutine

### Purpose

Returns the boolean entry for the specified identifier.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
bool tigetflag( ID)  
char *ID;
```

### Description

The **tigetflag** subroutine returns the boolean entry for the specified **termcap** identifier. This subroutine is provided for binary compatibility with applications that use the **termcap** file.

### Parameters

*ID*        Specifies the 2-character string that contains a **termcap** identifier.

### Return Values

The **tigetflag** subroutine returns the boolean entry for the specified **termcap** identifier. If *ID* is not found, on not a boolean, 0 is returned.

### Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tgetnum Subroutine

### Purpose

Returns the numeric entry for the specified **termcap** identifier.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
int tgetnum( ID)  
char *ID;
```

### Description

The **tgetnum** subroutine returns the numeric entry for the specified **termcap** identifier. This subroutine is provided for binary compatibility with applications that use the **termcap** file.

### Parameters

*ID* Specifies the 2-character string that contains a **termcap** identifier.

### Return Values

The **tgetnum** subroutine returns the numeric entry for the specified **termcap** identifier.

-1 Returned if the ID is not found or not numeric.

### Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tgetstr Subroutine

### Purpose

Returns the string entry for the specified **termcap** identifier.

## Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>
```

```
char *tgetstr( ID, Area)
char *ID, **Area;
```

## Description

The **tgetstr** subroutine returns the string entry for the specified **termcap** identifier. This subroutine is provided for binary compatibility with applications that use the **termcap** file.

## Parameters

*Area* Contains the string entry for the specified termcap identifier. This also is returned to the calling program.  
*ID* Specifies the 2-character string that contains the **termcap** identifier.

## Return Values

The **tgetstr** subroutine returns the string entry for the *ID* parameter, which is a 2-character string that contains a **termcap** identifier.

0 Returned if ID is not found or not a string capability.

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tgoto Subroutine

### Purpose

Duplicates the **tparm** subroutine.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
#include <term.h>
```

```
char *tgoto( Capability, Column, Row)
char *Capability;
int Column, Row;
```

## Description

The **tgoto** subroutine calls the **tparam** (“tparam Subroutine” on page 651) subroutine. This subroutine is provided for binary compatibility with applications that use the **termcap** file.

## Parameters

<i>Capability</i>	Specifies the termcap capability to apply the parameters to.
<i>Column</i>	Specifies which column to apply to the capability.
<i>Row</i>	Specifies which row to apply to the capability.

## Related Information

The **tparam** (“tparam Subroutine” on page 651) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tigetflag, tigetnum, tigetstr, or tparam Subroutine

### Purpose

Retrieves capabilities from the **terminfo** database.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <term.h>

int tigetflag(char *capname,);

int tigetnum(char *capname);

char *tigetstr(char *capname);

char *tparam(char *cap,
long p1, long p2, long p3,
long p4, long p5, long p6
long p7, long p8, long p9);
```

### Description

The **tigetflag**, **tigetnum**, and **tigetstr** subroutines obtain boolean, numeric, and string capabilities, respectively, from the selected record of the terminfo database. For each capability, the value to use as capname appears in the Capname column in the table in Section 6.1.3 on page 296.

The **tparam** subroutine takes as *cap* a string capability. If *cap* is parameterised (as described in Section A.1.2 on page 313), the **tparam** subroutine resolves the parameterisation. If the parameterised string refers to parameters *%p1* through *%p9*, then the **tparam** subroutine substitutes the values of *p1* through *p9*, respectively.

### Return Values

Upon successful completion, the **tigetflag**, **tigetnum**, and **tigetstr** subroutines return the specified capability. The **tigetflag** subroutine returns -1 if capname is not a boolean capability. The **tigetnum** subroutine returns -2 if capname is not a numeric capability. The **tigetstr** subroutine returns (char\*)-1 if capname is not a string capability.

Upon successful completion, the **tparm** subroutine returns *str* with parameterisation resolved. Otherwise, it returns a null pointer.

## Parameters

*\*capname*

*\*tparm*

*long p1*

*long p2*

*long p3*

*long p4*

*long p5*

*long p6*

*long p7*

*long p8*

*long p9*

## Examples

For the **tigetflag** subroutine:

To determine if erase overstrike is a defined boolean capability for the current terminal, use:

```
rc = tigetflag("eo");
```

For the **tigetnum** subroutine:

To determine if number of labels is a defined numeric capability for the current terminal, use:

```
rc = tigetnum("nlab");
```

For the **tigetstr** subroutine:

To determine if "turn on soft labels" is a defined string capability for the current terminal, do the following:

```
char *rc;  
rc = tigetstr("smln");
```

For the **tparm** subroutine:

1. To save the escape sequence used to home the cursor in the user-defined variable `home_sequence`, enter:

```
home_sequence = tparm(cursor_home);
```

2. To save the escape sequence used to move the cursor to the coordinates X=40, Y=18 in the user-defined variable `move_sequence`, enter:

```
move_sequence = tparm(cursor_address, 18, 40);
```

## Related Information

The **def\_prog\_mode** ("def\_prog\_mode, def\_shell\_mode, reset\_prog\_mode or reset\_shell\_mode Subroutine" on page 545), **tgetent** ("tgetent, tgetflag, tgetnum, tgetstr, or tgoto Subroutine" on page 640), and **putp** ("putp, tputs Subroutine" on page 606) subroutines.

Curses Overview for Programming, List of Curses Subroutines

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tigetnum Subroutine

### Purpose

Gets the value of terminal's numeric capability.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
#include <term.h>
```

```
tigetnum( CapName)
register char *CapName;
```

### Description

The **tigetnum** subroutine returns the value of terminal's numeric capability. Use this subroutine to get a capability for the current terminal. When successful, this subroutine returns the current value of the capability specified by the *CapName* parameter. Otherwise, if it is not a numeric value, this subroutine returns -2.

**Note:** The **tigetnum** subroutine is a low-level routine. Use this subroutine only if your application must deal directly with the terminfo database to handle certain terminal capabilities (for example, programming function keys).

### Return Values

Upon successful completion, the **tigetnum** subroutine returns the value of terminal's numeric capability.

-2     Indicates the value specified by the *CapName* parameter is not numeric.

### Parameters

*CapName*     Identifies the terminal capability to check for.

### Example

To determine if number of labels is a defined numeric capability for the current terminal, use:

```
rc = tigetnum("nlab");
```

### Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tigetstr Routine

### Purpose

Returns the value of a terminal's string capability.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
#include <term.h>
```

```
tigetstr( Capname)
register char *Capname;
```

### Description

The **tigetstr** subroutine returns the value of terminal's string capability. Use this subroutine to get a capability for the current terminal pointed to by **cur\_term**. When successful, this subroutine returns the current value of the capability specified by the *Capname* parameter. Otherwise, if it is not a string value, this subroutine returns (**char\***) -1.

**Note:** The **tigetstr** subroutine is a low-level routine. Use this subroutine only if your application must deal directly with the terminfo database to handle certain terminal capabilities (for example, programming function keys).

### Parameters

*Capname*      Identifies the terminal capability to check.

### Example

To determine if "turn on soft labels" is a defined string capability for the current terminal, do the following:

```
char *rc;
rc = tigetstr("smln");
```

### Return Values

Upon successful completion, the **tigetstr** subroutine returns the value of terminal's string capability.

(**char \***)-1      Indicates the value specified by the *Capname* parameter is not a string.

### Files

**/usr/include/curses.h**      Contains C language subroutines and define statements for curses.

### Related Information

List of Curses Subroutines, Curses Overview for Programming, Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## is\_linetouched, is\_wintouched, touchline, touchwin, untouchwin, or wtouchin Subroutine

### Purpose

Window refresh control functions.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
bool is_linetouched(WINDOW *win,  
int line);
```

```
bool is_wintouched(WINDOW *win);
```

```
int touchline(WINDOW *win,  
int start,  
int count);
```

```
int touchwin(WINDOW *win);
```

```
int untouchwin(WINDOW *win);
```

```
int wtouchln(WINDOW *win,  
int y,  
int n,  
int changed);
```

### Description

The **touchline** subroutine touches the specified window (that is, marks it as having changed more recently than the last refresh operation). The **touchline** subroutine only touches count lines, beginning with line start.

The **untouchwin** subroutine marks all lines in the window as unchanged since the last refresh operation.

Calling the **wtouchln** subroutine, if changed is 1, touches n lines in the specified window, starting at line y. If changed is 0, **wtouchln** marks such lines as unchanged since the last refresh operation.

The **is\_wintouchwin** subroutine determines whether the specified window is touched. The **is\_linetouched** subroutine determines whether line line of the specified window is touched.

### Parameters

*line*

*start*

*count*

*changed*

*y*

*n*

*\*win*

## Return Values

The `is_linetouched` and `is_wintouched` subroutines return TRUE if any of the specified lines, or the specified window, respectively, has been touched since the last refresh operation. Otherwise, they return FALSE.

Upon successful completion, the other subroutines return OK. Otherwise, they return ERR. Exceptions to this are noted in the preceding subroutine.

## Examples

For the `touchline` subroutine:

To set 10 lines for refresh starting from line 5 of the user-defined window `my_window`, use:

```
WINDOW *my_window;
touchline(my_window, 5, 10);
wrefresh(my_window);
```

This forces **curses** to disregard any optimization information it may have for lines 0-4 in `my_window`. **curses** assumes all characters in lines 0-4 have changed.

For the `touchwin` subroutine:

To refresh a user-defined parent window, `parent_window`, that has been edited through its subwindows, use:

```
WINDOW *parent_window;
touchwin(parent_window);

wrefresh(parent_window);
```

This forces **curses** to disregard any optimization information it may have for `my_window`. **curses** assumes all lines and columns have changed for `my_window`.

## Related Information

The **douupdate** (“`douupdate`, `refresh`, `wnoutrefresh`, or `wrefresh` Subroutines” on page 657) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Windows with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## touchoverlap Subroutine

### Purpose

Marks the overlap of two windows as changed and makes arrangements for their refresh.

### Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>
```

```
touchoverlap( Window1, Window2)  
WINDOW *Window1, Window2;
```

## Description

The **touchoverlap** subroutine marks the overlap of two windows as changed and makes arrangements for their refresh.

## Parameters

*Window1*      Specifies the first window as changed.  
*Window2*      Specifies the second window as changed.

## Examples

To mark the overlap of the two user-defined windows `my_window` and `my_new_window` as changed, enter:  
`touchoverlap(my_window, my_new_window);`

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Windows with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## touchwin Subroutine

### Purpose

Forces every character in a window's buffer to be refreshed at the next call to the **wrefresh** subroutine.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
touchwin( Window)  
WINDOW *Window;
```

### Description

The **touchwin** (“touchwin Subroutine”) subroutine forces every character in the specified window to be refreshed during the next call to the **refresh** or **wrefresh** subroutine. To force a specific range of lines to be refreshed, use the **touchline** (“is\_linetouched, is\_wintouched, touchline, touchwin, untouchwin, or wtouchin Subroutine” on page 648) subroutine.

The combined usage of the **touchwin** and **wrefresh** subroutines is helpful when dealing with subwindows or overlapping windows. When dealing with overlapping windows, it may become necessary to bring the back window to the front. A call to the **wrefresh** subroutine does not change the terminal because none of the characters in the window were changed. Calling the **touchwin** subroutine on the back window before the **wrefresh** subroutine redisplay the window on the terminal and, effectively, brings it to the front.

## Parameters

*Window* Specifies the window to be touched.

## Example

To refresh a user-defined parent window, `parent_window`, that has been edited through its subwindows, use:

```
WINDOW *parent_window;
touchwin(parent_window);

wrefresh(parent_window);
```

This forces **curses** to disregard any optimization information it may have for `my_window`. **curses** assumes all lines and columns have changed for `my_window`.

## Related Information

The **touchline** (“`is_linetouched`, `is_wintouched`, `touchline`, `touchwin`, `untouchwin`, or `wtouchin` Subroutine” on page 648) subroutine, **wrefresh** (“`refresh` or `wrefresh` Subroutine” on page 608) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Windows in the Curses Environment in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tparm Subroutine

### Purpose

Applies parameters (padding) to a terminal capability.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
char *tparm( TermCap, Parm1, Parm2, . . . Parm9)
char *TermCap;
int Parm1, Parm2, . . . Parm9;
```

### Description

The **tparm** subroutine applies parameters (padding) to a terminal capability.

**Note:** If the **tparm** subroutine is called with less than 10 parameters, then the **-D\_TPARM\_COMPAT** option should be used when compiling the program. Otherwise the compiler gives the following error.

```
1506-098 (E) Missing argument(s)
```

## Parameters

<i>Parm#</i>	Specifies the parameters (up to nine) to instantiate.
<i>TermCap</i>	Specifies the terminal capability to apply the parameters to. These terminal capabilities are defined in the <b>term.h</b> file.

## Return Values

The **tparam** subroutine returns the escape sequence specified by the *TermCap* parameter with the specified parameters applied. After the escape sequence is received, it can be output by a subroutine like the **tputs** (“tputs Subroutine”) subroutine.

## Examples

1. To save the escape sequence used to home the cursor in the user-defined variable `home_sequence`, enter:

```
home_sequence = tparm(cursor_home);
```

2. To save the escape sequence used to move the cursor to the coordinates X=40, Y=18 in the user-defined variable `move_sequence`, enter:

```
move_sequence = tparm(cursor_address, 18, 40);
```

## Related Information

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## tputs Subroutine

### Purpose

Outputs a string with padding information.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
#include <term.h>
```

```
tputs( String, LinesAffected, PutcLikeSub)
char *String;
int LinesAffected;
int (*PutcLikeSub) ();
```

### Description

The **tputs** subroutine outputs a string with padding information applied. String must be a terminfo string variable or the return value from **tparam**, **tgetstr**, **tigetstr**, or **tgoto** subroutines.

## Parameters

<i>LinesAffected</i>	Specifies the number of lines affected, or specifies 1 if not applicable.
<i>PutcLikeSub</i>	Specifies a <b>putchar</b> -like subroutine through which the characters are passed one at a time.
<i>String</i>	Specifies the string to which to add padding information.

## Examples

1. To output the clear screen sequence using the user-defined **putchar**-like subroutine **my\_putchar**, enter:  

```
int_my_putchar();  
tputs(clear_screen, 1 ,my_putchar);
```
2. To output the escape sequence used to move the cursor to the coordinates x=40, y=18 through the user-defined **putchar**-like subroutine **my\_putchar**, enter:  

```
int_my_putchar();  
tputs(tparm(cursor_address, 18, 40), 1, my_putchar);
```

## Related Information

The **tparm** (“tparm Subroutine” on page 651) subroutine.

Curses Overview for Programming, List of Curses Subroutines, Understanding Terminals with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## typeahead Subroutine

### Purpose

Controls checking for typeahead.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>  
  
int typeahead  
(int fildes);
```

### Description

The **typeahead** subroutine controls the detection of typeahead during a refresh, based on the value of *fildes*:

- If *fildes* is a valid file descriptor, the **typeahead** subroutine is enabled during refresh; Curses periodically checks *fildes* for input and aborts refresh if any character is available. (This is the initial setting, and the *typeahead* file descriptor corresponds to the input file associated with the screen created by the **initscr** or **newterm** subroutine.) The value of *fildes* need not be the file descriptor on which the refresh is occurring.
- If *fildes* is -1, Curses does not check for typeahead during refresh.

### Parameters

*fildes*

## Return Value

Upon successful completion, the **typeahead** subroutine returns OK. Otherwise, it returns ERR.

## Example

To turn typeahead checking on, enter:

```
typeahead(1);
```

## Related Information

The **douupdate** (“douupdate, refresh, wnoutrefresh, or wrefresh Subroutines” on page 657), **getch** (“getch, mvgetch, mvwgetch, or wgetch Subroutine” on page 561), and **initscr** (“initscr and newterm Subroutine” on page 577) subroutines.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Setting Video Attributes and Curses Options in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## unctrl Subroutine

### Purpose

Generates a printable representation of a character.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
```

```
char *unctrl  
(chtype c);
```

### Description

The **unctrl** subroutine generates a character string that is a printable representation of *c*. If *c* is a control character, it is converted to the ^X notation. If *c* contains rendition information, the effect is undefined.

### Parameters

*c*

### Return Values

Upon successful completion, the **unctrl** subroutine returns the generated string. Otherwise, it returns a null pointer.

### Examples

To display a printable representation of the newline character, enter:

```
char *new_line;
int my_character;
addstr ("Hit the enter key.");
my_character=getch();
new_line=unctrl (my_character);
printw (Newline=%s", new_line);
refresh();
```

This prints, "newline=^J".

## Related Information

The **keyname** ("keyname, key\_name Subroutine" on page 581) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## ungetch, unget\_wch Subroutine

### Purpose

Pushes a character onto the input queue.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>
int ungetch
(int ch);
int unget_wch
(const wchar_t wch);
```

### Description

The **ungetch** subroutine pushes the single-byte character *ch* onto the head of the input queue.

The **unget\_wch** subroutine pushes the wide character *wch* onto the head of the input queue.

One character of push-back is guaranteed. The result of successive calls without an intervening call to the **getch** or **get\_wch** subroutine are unspecified.

### Parameters

*ch*  
*wch*

## Examples

To force the key KEY\_ENTER back into the queue, use:

```
ungetch(KEY_ENTER);
```

## Related Information

The **getch** or **wgetch** (“getch, mvgetch, mvwgetch, or wgetch Subroutine” on page 561) subroutine.

Curses Overview for Programming and List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Characters with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## vidattr, vid\_attr, vidputs, or vid\_puts Subroutine

### Purpose

Outputs attributes to the terminal.

### Library

Curses Library (**libcurses.a**)

### Syntax

```
#include <curses.h>

int vidattr
(chtype attr);

int vid_attr
(attr_t attr,
short color_pair_number,
void *opt);

int vidputs
(chtype attr,
int (*putfunc)(int));

int vid_puts
(attr_t attr,
short color_pair_number,
void *opt,
int (*putfunc)(int));
```

### Description

These subroutines output commands to a terminal that changes the terminal's attributes.

If the **terminfo** database indicates that the terminal in use can display characters in the rendition specified by *attr*, then the **vidattr** subroutine outputs one or more commands to request that the terminal display subsequent characters in that rendition. The subroutine outputs by calling the **putchar** subroutine. The **vidattr** subroutine neither relies on nor updates the model that Curses maintains of the prior rendition mode.

The **vidputs** subroutine computes the same terminal output string that **vidattr** does, based on *attr*, but the **vidputs** subroutine outputs by calling the user-supplied subroutine **putfunc**. The **vid\_attr** and **vid\_puts** subroutines correspond to **vidattr** and **vidputs** respectively, but take a set of arguments, one of type *attr\_t* for the attributes, *short* for the color pair number and a *void \**, and thus support the attribute constants with the *WA\_* prefix.

The *opts* argument is reserved for definition in a future edition of this document. Currently, the application must provide a null pointer as *opts*.

The user-supplied **putfunc** subroutine (which can be specified as an argument to either **vidputs** or **vid\_puts** is either **putchar** or some other subroutine with the same prototype. Both the **vidputs** and the **vid\_puts** subroutines ignore the return value of **putfunc**.

## Parameters

*att*  
*color\_pair\_number*  
*\*opt*  
*\*putfunc*

## Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

## Examples

1. To output the string that puts the terminal in its best standout mode through the **putchar** subroutine, enter  

```
vidattr(A_STANDOUT);
```
2. To output the string that puts the terminal in its best standout mode through the **putchar**-like subroutine **my\_putc**, enter  

```
int (*my_putc) ();  
vidputs(A_STANDOUT, my_putc);
```

## Related Information

The **douupdate** (“**douupdate**, **refresh**, **wnoutrefresh**, or **wrefresh** Subroutines”), **is\_linetouched** (“**is\_linetouched**, **is\_wintouched**, **touchline**, **touchwin**, **untouchwin**, or **wtouchin** Subroutine” on page 648), **putchar**, **putwchar** and **tigetflag** (“**tigetflag**, **tigetnum**, **tigetstr**, or **tparm** Subroutine” on page 644) subroutines.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Setting Video Attributes and Curses Options in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## douupdate, refresh, wnoutrefresh, or wrefresh Subroutines

### Purpose

Refreshes windows and lines.

### Library

Curses Library (**libcurses.a**)

## Syntax

```
#include <curses.h>

int doupdate(void);

int refresh(void);

int wnoutrefresh(WINDOW *win);

int wrefresh(WINDOW *win);
```

## Description

The **refresh** and **wrefresh** subroutines refresh the current or specified window. The subroutines position the terminal's cursor at the cursor position of the window, except that, if the leaveok mode has been enabled, they may leave the cursor at an arbitrary position.

The **wnoutrefresh** subroutine determines which parts of the terminal may need updating.

The **doupdate** subroutine sends to the terminal the commands to perform any required changes.

## Parameters

*\*win* Specifies the window to be refreshed.

## Return Values

Upon successful completion, these subroutines return OK. Otherwise, they return ERR.

## Examples

For the **doupdate** or **wnoutrefresh** subroutine:

To update the user-defined windows `my_window1` and `my_window2`, enter:

```
WINDOW *my_window1, my_window2;
wnoutrefresh(my_window1);
wnoutrefresh(my_window2);
doupdate();
```

For the **refresh** or **wrefresh** subroutine:

1. To update the terminal's display and the current screen structure to reflect changes made to the standard screen structure, use:  
`refresh();`
2. To update the terminal and the current screen structure to reflect changes made to a user-defined window called `my_window`, use:  
`WINDOW *my_window;`  
`wrefresh(my_window);`
3. To restore the terminal to its state at the last refresh, use:  
`wrefresh(curscr);`

This subroutine is useful if the terminal becomes garbled for any reason.

## Related Information

The **clearok** (“clearok, idlok, leaveok, scrollok, setscreg or wsetscreg Subroutine” on page 537) subroutine.

Curses Overview for Programming in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

List of Curses Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

Manipulating Window Data with Curses in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.



---

## Chapter 3. FORTRAN Basic Linear Algebra Subroutines (BLAS)

---

### SDOT or DDOT Function

#### Purpose

Returns the dot product of two vectors.

#### Library

BLAS Library (**libblas.a**)

#### FORTRAN Syntax

```
REAL FUNCTION SDOT(N, X, INCX, Y, INCY)
INTEGER INCX, INCY, N
REAL X(*), Y(*)

DOUBLE PRECISION FUNCTION DDOT(N, X, INCX, Y, INCY)
INTEGER INCX, INCY, N
DOUBLE PRECISION X(*), Y(*)
```

#### Description

The **SDOT** or **DDOT** function returns the dot product of vectors *X* and *Y*.

#### Parameters

*N* On entry, *N* specifies the number of elements in *X* and *Y*; unchanged on exit.  
*X* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; unchanged on exit.  
*INCX* On entry, *INCX* specifies the increment for the elements of *X*; unchanged on exit.  
*Y* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCY}))$ ; unchanged on exit.  
*INCY* On entry, *INCY* specifies the increment for the elements of *Y*; unchanged on exit.

#### Error Codes

For values of *N*  $\leq 0$ , a value of 0 is returned.

---

### CDOTC or ZDOTC Function

#### Purpose

Returns the complex dot product of two vectors, conjugating the first.

#### Library

BLAS Library (**libblas.a**)

#### FORTRAN Syntax

```
COMPLEX FUNCTION CDOTC(N, X, INCX, Y, INCY)
INTEGER INCX, INCY, N
COMPLEX X(*), Y(*)
```

**DOUBLE COMPLEX FUNCTION ZDOTC**(*N*, *X*, *INCX*, *Y*, *INCY*)  
**INTEGER** *INCX*, *INCY*, *N*  
**COMPLEX\*16** *X*(\*), *Y*(\*)

## Description

The **CDOTC** or **ZDOTC** function returns the complex dot product of two vectors, conjugating the first.

## Parameters

*N* On entry, *N* specifies the number of elements in *X* and *Y*; unchanged on exit.  
*X* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; unchanged on exit.  
*INCX* On entry, *INCX* specifies the increment for the elements of *X*; unchanged on exit.  
*Y* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCY}))$ ; unchanged on exit.  
*INCY* On entry, *INCY* specifies the increment for the elements of *Y*; unchanged on exit.

## Error Codes

For values of *N*  $\leq 0$ , a value of 0 is returned.

---

## CDOTU or ZDOTU Function

### Purpose

Returns the complex dot product of two vectors.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

**COMPLEX FUNCTION CDOTU**(*N*, *X*, *INCX*, *Y*, *INCY*)  
**INTEGER** *INCX*, *INCY*, *N*  
**COMPLEX** *X*(\*), *Y*(\*)  
**DOUBLE COMPLEX FUNCTION ZDOTU**(*N*, *X*, *INCX*, *Y*, *INCY*)  
**INTEGER** *INCX*, *INCY*, *N*  
**COMPLEX\*16** *X*(\*), *Y*(\*)

## Description

The **CDOTU** or **ZDOTU** function returns the complex dot product of two vectors.

## Parameters

*N* On entry, *N* specifies the number of elements in *X* and *Y*; unchanged on exit.  
*X* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; unchanged on exit.  
*INCX* On entry, *INCX* specifies the increment for the elements of *X*; unchanged on exit.  
*Y* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCY}))$ ; unchanged on exit.  
*INCY* On entry, *INCY* specifies the increment for the elements of *Y*; unchanged on exit.

## Error Codes

For values of *N*  $\leq 0$ , a value of 0 is returned.

---

## SAXPY, DAXPY, CAXPY, or ZAXPY Subroutine

### Purpose

Computes a constant times a vector plus a vector.

### Library

BLAS Library (**libblas.a**)

### FORTTRAN Syntax

```
SUBROUTINE SAXPY(N,A,X,INCX,Y,INCY)
```

```
INTEGER INCX, INCY, N
```

```
REAL A
```

```
REAL X(*), Y(*)
```

```
SUBROUTINE DAXPY(N,A,X,INCX,Y,INCY)
```

```
INTEGER INCX, INCY, N
```

```
DOUBLE PRECISION A
```

```
DOUBLE PRECISION X(*), Y(*)
```

```
SUBROUTINE CAXPY(N,A,X,INCX,Y,INCY)
```

```
INTEGER INCX, INCY, N
```

```
COMPLEX A
```

```
COMPLEX X(*), Y(*)
```

```
SUBROUTINE ZAXPY(N,A,X,INCX,Y,INCY)
```

```
INTEGER INCX, INCY, N
```

```
COMPLEX*16 A
```

```
COMPLEX*16 X(*), Y(*)
```

### Description

The **SAXPY**, **DAXPY**, **CAXPY**, or **ZAXPY** subroutine computes a constant times a vector plus a vector:

$$Y = A * X + Y$$

### Parameters

- N* On entry, *N* specifies the number of elements in *X* and *Y*; unchanged on exit.
- A* On entry, *A* contains a constant to be multiplied by the *X* vector; unchanged on exit.
- X* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; unchanged on exit.
- INCX* On entry, *INCX* specifies the increment for the elements of *X*; unchanged on exit.
- Y* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCY}))$ ; the result is returned in vector *Y*.
- INCY* On entry, *INCY* specifies the increment for the elements of *Y*; unchanged on exit.

### Error Codes

If *SA* = 0 or *N* <= 0, the subroutine returns immediately.

---

## SROTG, DROTG, CROTG, or ZROTG Subroutine

### Purpose

Constructs Givens plane rotation.

### Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

**SUBROUTINE SROTG**(*A,B,C,S*)

**REAL** *A, B, C, S*

**SUBROUTINE DROTG**(*A,B,C,S*)

**DOUBLE PRECISION** *A,B,C,S*

**SUBROUTINE CROTG**(*A,B,C,S*)

**REAL** *C*

**COMPLEX** *A,B,S*

**SUBROUTINE ZROTG**(*A,B,C,S*)

**DOUBLE PRECISION** *C*

**COMPLEX\*16** *A,B,S*

## Description

Given vectors *A* and *B*, the **SROTG**, **DROTG**, **CROTG**, or **ZROTG** subroutine computes:

$$a = \frac{A}{|A| + |B|}, \quad b = \frac{B}{|A| + |B|}$$

$$r = \begin{cases} a & \text{if } |A| > |B| \\ b & \text{if } |B| \geq |A| \end{cases} \quad r = r \sqrt{a^2 + b^2}$$

$$C = \begin{cases} A/r & \text{if } r \text{ not} = 0 \\ 1 & \text{if } r = 0 \end{cases} \quad S = \begin{cases} B/r & \text{if } r \text{ not} = 0 \\ 0 & \text{if } r = 0 \end{cases}$$

The numbers *C*, *S*, and *r* then satisfy the matrix equation:

$$\begin{bmatrix} C & S \\ -S & C \end{bmatrix} \cdot \begin{bmatrix} A \\ B \end{bmatrix} = \begin{bmatrix} r \\ 0 \end{bmatrix}$$

The subroutines also compute:

$$z = \begin{cases} S & \text{if } |A| > |B|, \\ 1/C & \text{if } |B| \geq |A| \text{ and } C \text{ not} = 0, \\ 1 & \text{if } C = 0. \end{cases}$$

The subroutines return *r* overwriting *A* and *z* overwriting *B*, as well as returning *C* and *S*.

## Parameters

- A* On entry, contains a scalar constant; on exit, contains the value *r*.
- B* On entry, contains a scalar constant; on exit, contains the value *z*.
- C* Can contain any value on entry; the value *C* returned on exit.
- S* Can contain any value on entry; the value *S* returned on exit.

---

## SROT, DROT, CSROT, or ZDROT Subroutine

### Purpose

Applies a plane rotation.

### Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

```
SUBROUTINE SROT(N,X,INCX,Y,INCY,C,S)
INTEGER INCX, INCY, N
REAL C, S
REAL X(*), Y(*)
SUBROUTINE DROT(N,X,INCX,Y,INCY,C,S)
INTEGER INCX, INCY, N
DOUBLE PRECISION C, S
DOUBLE PRECISION X(*), Y(*)
SUBROUTINE CSROT(N,X,INCX,Y,INCY,C,S)
INTEGER INCX, INCY, N
REAL C, S
COMPLEX X(*), Y(*)
SUBROUTINE ZDROT(N,X,INCX,Y,INCY,C,S)
INTEGER INCX, INCY, N
DOUBLE PRECISION C, S
COMPLEX*16 X(*), Y(*)
```

## Description

The **SROT**, **DROT**, **CSROT**, or **ZDROT** subroutine computes:

$$\begin{array}{|c|} \hline X \\ \hline i \\ \hline Y \\ \hline i \\ \hline \end{array} := \begin{array}{|cc|} \hline C & S \\ \hline -S & C \\ \hline \end{array} \cdot \begin{array}{|c|} \hline X \\ \hline i \\ \hline Y \\ \hline i \\ \hline \end{array} \quad \text{for } i = 1, \dots, N.$$

The subroutines return the modified  $X$  and  $Y$ .

## Parameters

- $N$  On entry,  $N$  specifies the number of elements in  $X$  and  $Y$ ; unchanged on exit.
- $X$  Vector of dimension at least  $(1 + (N-1) * \text{abs}(INCX))$ ; unchanged on exit.
- $INCX$  On entry,  $INCX$  specifies the increment for the elements of  $X$ ; unchanged on exit.
- $Y$  Vector of dimension at least  $(1 + (N-1) * \text{abs}(INCY))$ ; modified on exit.
- $INCY$  On entry,  $INCY$  specifies the increment for the elements of  $Y$ ; unchanged on exit.
- $C$  Scalar constant; unchanged on exit.
- $S$  Scalar constant; unchanged on exit.

## Error Codes

If  $N \leq 0$ , or if  $C = 1$  and  $S = 0$ , the subroutines return immediately.

---

## SCOPY, DCOPY, CCOPY, or ZCOPY Subroutine

### Purpose

Copies vector  $X$  to  $Y$ .

### Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

```
SUBROUTINE SCOPY(N,X,INCX,Y,INCY)
```

```
INTEGER INCX, INCY, N
```

```
REAL X(*), Y(*)
```

```
SUBROUTINE DCOPY(N,X,INCX,Y,INCY)
```

```
INTEGER INCX, INCY, N
```

```
DOUBLE PRECISION X(*), Y(*)
```

```
SUBROUTINE CCOPY(N,X,INCX,Y,INCY)
```

```
INTEGER INCX, INCY, N
```

```
COMPLEX X(*), Y(*)
```

```
SUBROUTINE ZCOPY(N,X,INCX,Y,INCY)
```

```
INTEGER INCX, INCY, N
```

```
COMPLEX*16 X(*), Y(*)
```

## Description

The **SCOPY**, **DCOPY**, **CCOPY**, or **ZCOPY** subroutine copies vector *X* to vector *Y*.

## Parameters

- N* On entry, *N* specifies the number of elements in *X* and *Y*; unchanged on exit.
- X* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; unchanged on exit.
- INCX* On entry, *INCX* specifies the increment for the elements of *X*; unchanged on exit.
- Y* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCY}))$  or greater; can contain any values on entry; on exit, contains the same values as *X*.
- INCY* On entry, *INCY* specifies the increment for the elements of *Y*; unchanged on exit.

## Error Codes

For values of *N*  $\leq 0$ , the subroutines return immediately.

---

## SSWAP, DSWAP, CSWAP, or ZSWAP Subroutine

### Purpose

Interchanges vectors *X* and *Y*.

### Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

```
SUBROUTINE SSWAP(N,X,INCX,Y,INCY)
```

```
INTEGER INCX, INCY, N
```

```
REAL X(*), Y(*)
```

```
SUBROUTINE DSWAP(N,X,INCX,Y,INCY)
```

```
INTEGER INCX, INCY, N
```

```
DOUBLE PRECISION X(*), Y(*)
```

```
SUBROUTINE CSWAP(N,X,INCX,Y,INCY)
```

```
INTEGER INCX, INCY, N
```

```
COMPLEX X(*), Y(*)
```

```
SUBROUTINE ZSWAP(N,X,INCX,Y,INCY)
```

```
INTEGER INCX, INCY, N
```

```
COMPLEX*16 X(*), Y(*)
```

## Description

The **SSWAP**, **DSWAP**, **CSWAP**, or **ZSWAP** subroutine interchanges vector  $X$  and vector  $Y$ .

## Parameters

$N$  On entry,  $N$  specifies the number of elements in  $X$  and  $Y$ ; unchanged on exit.  
 $X$  Vector of dimension at least  $(1 + (N-1) * \text{abs}(INCX))$ ; on exit, contains the elements of vector  $Y$ .  
 $INCX$  On entry,  $INCX$  specifies the increment for the elements of  $X$ ; unchanged on exit.  
 $Y$  Vector of dimension at least  $(1 + (N-1) * \text{abs}(INCY))$ ; on exit, contains the elements of vector  $X$ .  
 $INCY$  On entry,  $INCY$  specifies the increment for the elements of  $Y$ ; unchanged on exit.

## Error Codes

For values of  $N \leq 0$ , the subroutines return immediately.

---

## SNRM2, DNRM2, SCNRM2, or DZNRM2 Function

### Purpose

Computes the Euclidean length of the  $N$ -vector stored in  $X()$  with storage increment  $INCX$ .

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
REAL FUNCTION SNRM2( $N,X,INCX$ )  
INTEGER  $INCX,N$   
REAL  $X(*)$   
  
DOUBLE PRECISION FUNCTION DNRM2( $N,X,INCX$ )  
INTEGER  $INCX,N$   
DOUBLE PRECISION  $X(*)$   
  
REAL FUNCTION SCNRM2( $N,X,INCX$ )  
INTEGER  $INCX,N$   
COMPLEX  $X(*)$   
  
DOUBLE PRECISION FUNCTION DZNRM2( $N,X,INCX$ )  
INTEGER  $INCX,N$   
COMPLEX*16  $X(*)$ 
```

### Description

The **SNRM2**, **DNRM2**, **SCNRM2**, or **DZNRM2** function returns the Euclidean norm of the  $N$ -vector stored in  $X()$  with storage increment  $INCX$ .

### Parameters

$N$  On entry,  $N$  specifies the number of elements in  $X$  and  $Y$ ; unchanged on exit.  
 $X$  Vector of dimension at least  $(1 + (N-1) * \text{abs}(INCX))$ ; unchanged on exit.  
 $INCX$  On entry,  $INCX$  specifies the increment for the elements of  $X$ ;  $INCX$  must be greater than 0; unchanged on exit.

### Error Codes

For values of  $N \leq 0$ , a value of 0 is returned.

---

## SASUM, DASUM, SCASUM, or DZASUM Function

### Purpose

Returns the sum of absolute values of vector components.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
REAL FUNCTION SASUM(N,X,INCX)
INTEGER INCX,N
REAL X(*)
DOUBLE PRECISION FUNCTION DASUM(N,X,INCX)
INTEGER INCX,N
DOUBLE PRECISION X(*)
REAL FUNCTION SCASUM(N,X,INCX)
INTEGER INCX,N
COMPLEX X(*)
DOUBLE PRECISION FUNCTION DZASUM(N,X,INCX)
INTEGER INCX,N
COMPLEX*16 X(*)
```

### Description

The **SASUM**, **DASUM**, **SCASUM**, or **DZASUM** function returns the sum of absolute values of vector components.

### Parameters

*N* On entry, *N* specifies the number of elements in *X* and *Y*; unchanged on exit.  
*X* Vector of dimension at least  $(1 + (N-1) * \text{abs}(INCX))$ ; unchanged on exit.  
*INCX* On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must be greater than 0; unchanged on exit.

### Error Codes

For values of *N*  $\leq 0$ , a value of 0 is returned.

---

## SSCAL, DSCAL, CSSCAL, CSCAL, ZDSCAL, or ZSCAL Subroutine

### Purpose

Scales a vector by a constant.

### Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

```
SUBROUTINE SSCAL(N,A,X,INCX)
INTEGER INCX, N
REAL A
REAL X(*)
SUBROUTINE DSCAL(N,A,X,INCX)
INTEGER INCX,N
DOUBLE PRECISION A
DOUBLE PRECISION X(*)
SUBROUTINE CSSCAL(N,A,X,INCX)
INTEGER INCX,N
REAL A
COMPLEX X(*)
SUBROUTINE CSCAL
INTEGER INCX,N
COMPLEX A
COMPLEX X(*)
SUBROUTINE ZDSCAL
INTEGER INCX,N
DOUBLE PRECISION A
COMPLEX*16 X(*)
SUBROUTINE ZSCAL(
INTEGER INCX,N
COMPLEX*16 A
COMPLEX*16 X(*)
```

## Description

The **SSCAL**, **DSCAL**, **CSSCAL**, **CSCAL**, **ZDSCAL**, or **ZSCAL** subroutine scales a vector by a constant:

$X := X * A$

## Parameters

*N* On entry, *N* specifies the number of elements in *X* and *Y*; unchanged on exit.  
*A* Scaling constant; unchanged on exit.  
*X* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; on exit, contains the scaled vector.  
*INCX* On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must be greater than 0; unchanged on exit.

## Error Codes

For values of *N*  $\leq 0$ , the subroutines return immediately.

---

## ISAMAX, IDAMAX, ICAMAX, or IZAMAX Function

### Purpose

Finds the index of element having maximum absolute value.

### Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

INTEGER FUNCTION ISAMAX(*N,X,INCX*)

INTEGER *INCX, N*

REAL *X*(\*)

INTEGER FUNCTION IDAMAX(*N,X,INCX*)

INTEGER *INCX, N*

DOUBLE PRECISION *X*(\*)

INTEGER FUNCTION ICAMAX(*N,X,INCX*)

INTEGER *INCX, N*

COMPLEX *X*(\*)

INTEGER FUNCTION IZAMAX(*N,X,INCX*)

INTEGER *INCX, N*

COMPLEX\*16 *X*(\*)

## Description

The **ISAMAX**, **IDAMAX**, **ICAMAX**, or **IZAMAX** function returns the index of element having maximum absolute value.

## Parameters

- N* On entry, *N* specifies the number of elements in *X* and *Y*; unchanged on exit.  
*X* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; unchanged on exit.  
*INCX* On entry, *INCX* specifies the increment for the elements of *X*; unchanged on exit.

## Error Codes

For values of *N*  $\leq 0$ , a value of 0 is returned.

---

## SDSDOT Function

### Purpose

Returns the dot product of two vectors plus a constant.

### Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

REAL FUNCTION SDSDOT(*N,B,X,INCX,Y,INCY*)

INTEGER *N, INCX, INCY*

REAL *B, X*(\*), *Y*(\*)

### Purpose

The **SDSDOT** function computes the sum of constant *B* and dot product of vectors *X* and *Y*.

**Note:** Computation is performed in double precision.

### Parameters

- N* On entry, *N* specifies the number of elements in *X* and *Y*; unchanged on exit.  
*B* Scalar; unchanged on exit.  
*X* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; unchanged on exit.

*INCX* On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must be greater than zero; unchanged on exit.

*Y* Vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; unchanged on exit.

*INCY* On entry, *INCY* specifies the increment for the elements of *Y*; *INCY* must be greater than 0; unchanged on exit.

## Error Codes

For values of  $N \leq 0$ , the subroutine returns immediately.

---

## SROTM or DROTM Subroutine

### Purpose

Applies the modified Givens transformation.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

**SUBROUTINE SROTM**(*N,X,INCX,Y,INCY,PARAM*)

**INTEGER** *N, INCX, INCY*

**REAL** *X(\*), Y(\*), PARAM(5)*

**SUBROUTINE DROTM**(*N,X,INCX,Y,INCY,PARAM*)

**INTEGER** *N,INCX,INCY*

**DOUBLE PRECISION** *X(\*),Y(\*),PARAM(5)*

### Description

Let *H* denote the modified Givens transformation defined by the parameter array *PARAM*. The **SROTM** or **DROTM** subroutine computes:

$$\begin{bmatrix} x \\ y \end{bmatrix} := H * \begin{bmatrix} x \\ y \end{bmatrix}$$

where *H* is a 2 x 2 matrix with the components defined by the elements of the array *PARAM* as follows:

```

if PARAM(1) == 0.0
  H(1,1) = H(2,2) = 1.0
  H(2,1) = PARAM(3)
  H(1,2) = PARAM(4)
if PARAM(1) == 1.0
  H(1,2) = H(2,1) = -1.0
  H(1,1) = PARAM(2)
  H(2,2) = PARAM(5)
if PARAM(1) == -1.0
  H(1,1) = PARAM(2)
  H(2,1) = PARAM(3)
  H(1,2) = PARAM(4)
  H(2,2) = PARAM(5)
if PARAM(1) == -2.0
  H = I (Identity matrix)

```

If  $N \leq 0$  or *H* is an identity matrix, the subroutines return immediately.

## Parameters

<i>N</i>	On entry, <i>N</i> specifies the number of elements in <i>X</i> and <i>Y</i> ; unchanged on exit.
<i>X</i>	Vector of dimension at least $(1 + (N-1) * \text{abs}(INCX))$ ; on exit, modified as described above.
<i>INCX</i>	On entry, <i>INCX</i> specifies the increment for the elements of <i>X</i> ; <i>INCX</i> must be greater than 0; unchanged on exit.
<i>Y</i>	Vector of dimension at least $(1 + (N-1) * \text{abs}(INCY))$ ; on exit, modified as described above.
<i>INCY</i>	On entry, <i>INCY</i> specifies the increment for the elements of <i>Y</i> ; <i>INCY</i> must be greater than 0; unchanged on exit.
<i>PARAM</i>	Vector of dimension (5); on entry, must be set as described above. Specifically, <i>PARAM</i> (1) is a flag and must have value of either 0.0, -1.0, 1.0, or 2.0; unchanged on exit.

## Related information

The **SROTMG** or **DROTMG** (“SROTMG or DROTMG Subroutine”) subroutine builds the *PARAM* array prior to use by the **SROTM** or **DROTM** subroutine.

---

## SROTMG or DROTMG Subroutine

### Purpose

Constructs a modified Givens transformation.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE SROTMG(D1,D2,X1,X2,PARAM)
```

```
REAL D1, D2, X1, X2, PARAM(5)
```

```
SUBROUTINE DROTMG(D1,D2,X1,X2,PARAM)
```

```
DOUBLE PRECISION D1,D2,X1,X2,PARAM(5)
```

### Description

The **SROTMG** or **DROTMG** subroutine constructs a modified Givens transformation. The input quantities *D1*, *D2*, *X1*, and *X2* define a 2-vector in partitioned form:

$$\begin{array}{|c|} \hline a1 \\ \hline a2 \\ \hline \end{array} = \begin{array}{|cc|} \hline \text{sqrt}(D1) & 0 \\ \hline 0 & \text{sqrt}(D2) \\ \hline \end{array} \begin{array}{|c|} \hline X1 \\ \hline X2 \\ \hline \end{array}$$

The subroutines determine the modified Givens rotation matrix *H* that transforms *X2* and, thus, *a2* to 0. A representation of this matrix is stored in the array *PARAM* as follows:

Case 1: *PARAM*(1) = 1.0

*PARAM*(2) = *H*(1,1)

*PARAM*(5) = *H*(2,2)

Case 2: *PARAM*(1) = 0.0

*PARAM*(3) = *H*(2,1)

*PARAM*(4) = *H*(1,2)

Case 3: *PARAM*(1) = -1.0

*H*(1,1) = *PARAM*(2)

*H*(2,1) = *PARAM*(3)

*H*(1,2) = *PARAM*(4)

*H*(2,2) = *PARAM*(5)

Case 4:  $PARAM(1) = -2.0$   
 $H = I$  (Identity matrix)

**Note:** Locations in *PARAM* not listed are left unchanged.

## Parameters

*D1* Nonnegative scalar; modified on exit to reflect the results of the transformation.  
*D2* Scalar; can be negative on entry; modified on exit to reflect the results of the transformation.  
*X1* Scalar; modified on exit to reflect the results of the transformation.  
*X2* Scalar; unchanged on exit.  
*PARAM* Vector of dimension (5); values on entry are unused; modified on exit as described above.

## Related Information

The **SROT**M and **DROT**M (“SROTM or DROTM Subroutine” on page 671) subroutines apply the Modified Givens Transformation.

---

## SGEMV, DGEMV, CGEMV, or ZGEMV Subroutine

### Purpose

Performs matrix-vector operation with general matrices.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

**SUBROUTINE SGEMV**(*TRANS, M, N, ALPHA, A, LDA, X, INCX, BETA, Y, INCY*)

**REAL** *ALPHA, BETA*

**INTEGER** *INCX, INCY, LDA, M, N*

**CHARACTER\*1** *TRANS*

**REAL** *A(LDA,\*), X(\*), Y(\*)*

**SUBROUTINE DGEMV**(*TRANS, M, N, ALPHA, A, LDA, X, INCX, BETA, Y, INCY*)

**DOUBLE PRECISION** *ALPHA, BETA*

**INTEGER** *INCX, INCY, LDA, M, N*

**CHARACTER\*1** *TRANS*

**DOUBLE PRECISION** *A(LDA,\*), X(\*), Y(\*)*

**SUBROUTINE CGEMV**(*TRANS, M, N, ALPHA, A, LDA, X, INCX, BETA, Y, INCY*)

**COMPLEX** *ALPHA, BETA*

**INTEGER** *INCX, INCY, LDA, M, N*

**CHARACTER\*1** *TRANS*

**COMPLEX** *A(LDA,\*), X(\*), Y(\*)*

**SUBROUTINE ZGEMV**(*TRANS, M, N, ALPHA, A, LDA, X, INCX, BETA, Y, INCY*)

**COMPLEX\*16** *ALPHA, BETA*

**INTEGER** *INCX, INCY, LDA, M, N*

**CHARACTER\*1** *TRANS*

**COMPLEX\*16** *A(LDA,\*), X(\*), Y(\*)*

## Description

The **SGEMV**, **DGEMV**, **CGEMV**, or **ZGEMV** subroutine performs one of the following matrix-vector operations:

$$y := \alpha * A * x + \beta * y$$

OR

$$y := \alpha * A^T * x + \beta * y$$

where  $\alpha$  and  $\beta$  are scalars,  $x$  and  $y$  are vectors, and  $A$  is an  $M$  by  $N$  matrix.

## Parameters

**TRANS** On entry, **TRANS** specifies the operation to be performed as follows:

**TRANS = 'N' or 'n'**

$$y := \alpha * A * x + \beta * y$$

**TRANS = 'T' or 't'**

$$y := \alpha * A^T * x + \beta * y$$

**TRANS = 'C' or 'c'**

$$y := \alpha * A^T * x + \beta * y$$

Unchanged on exit.

**M** On entry, **M** specifies the number of rows of the matrix **A**; **M** must be at least 0; unchanged on exit.

**N** On entry, **N** specifies the number of columns of the matrix **A**; **N** must be at least 0; unchanged on exit.

**ALPHA** On entry, **ALPHA** specifies the scalar  $\alpha$ ; unchanged on exit.

**A** An array of dimension ( **LDA**, **N** ); on entry, the leading  $M$  by  $N$  part of the array **A** must contain the matrix of coefficients; unchanged on exit.

**LDA** On entry, **LDA** specifies the first dimension of **A** as declared in the calling (sub) program; **LDA** must be at least  $\max(1, M)$ ; unchanged on exit.

**X** A vector of dimension at least  $(1 + (N-1) * \text{abs}( \text{INCX} ))$  when **TRANS** = 'N' or 'n', otherwise, at least  $(1 + (M-1) * \text{abs}( \text{INCX} ))$ ; on entry, the incremented array **X** must contain the vector  $x$ ; unchanged on exit.

**INCX** On entry, **INCX** specifies the increment for the elements of **X**; **INCX** must not be 0; unchanged on exit.

**BETA** On entry, **BETA** specifies the scalar  $\beta$ ; when **BETA** is supplied as 0, **Y** need not be set on input; unchanged on exit.

**Y** A vector of dimension at least  $1 + (M-1) * \text{abs}( \text{INCY} )$  when **TRANS** = 'N' or 'n', otherwise at least  $(1 + (N-1) * \text{abs}( \text{INCY} ))$ ; on entry, with **BETA** nonzero, the incremented array **Y** must contain the vector  $y$ ; on exit, **Y** is overwritten by the updated vector  $y$ .

**INCY** On entry, **INCY** specifies the increment for the elements of **Y**; **INCY** must not be 0; unchanged on exit.

---

## SGBMV, DGBMV, CGBMV, or ZGBMV Subroutine

### Purpose

Performs matrix-vector operations with general banded matrices.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE SGBMV(TRANS, M, N, KL, KU, ALPHA, A, LDA,  
X, INCX, BETA, Y, INCY)
```

```
REAL ALPHA, BETA
```

```

INTEGER INCX, INCY, KL, KU, LDA, M, N
CHARACTER*1 TRANS
REAL A(LDA,*), X(*), Y(*)
SUBROUTINE DGBMV(TRANS, M, N, KL, KU, ALPHA, A, LDA,
X, INCX, BETA, Y, INCY)
DOUBLE PRECISION ALPHA,BETA
INTEGER INCX, INCY, KL, KU, LDA, M, N
CHARACTER*1 TRANS
DOUBLE PRECISION A(LDA,*), X(*), Y(*)
SUBROUTINE CGBMV(TRANS, M, N, KL, KU, ALPHA, A, LDA,
X, INCX, BETA, Y, INCY)
COMPLEX ALPHA,BETA
INTEGER INCX, INCY, KL, KU, LDA, M, N
CHARACTER*1 TRANS
COMPLEX A(LDA,*), X(*), Y(*)
SUBROUTINE ZGBMV(TRANS, M, N, KL, KU, ALPHA, A, LDA,
X, INCX, BETA, Y, INCY)
COMPLEX*16 ALPHA,BETA
INTEGER INCX, INCY, KL, KU, LDA, M, N
CHARACTER*1 TRANS
COMPLEX*16 A(LDA,*), X(*), Y(*)

```

## Description

The **SGBMV**, **DGBMV**, **CGBMV**, or **ZGBMV** subroutine performs one of the following matrix-vector operations:

$y := \alpha * A * x + \beta * y$

OR

$y := \alpha * A' * x + \beta * y$

where alpha and beta are scalars, x and y are vectors and A is an M by N band matrix, with KL subdiagonals and KU superdiagonals.

## Parameters

**TRANS** On entry, *TRANS* specifies the operation to be performed as follows:

*TRANS* = 'N' or 'n'

$y := \alpha * A * x + \beta * y$

*TRANS* = 'T' or 't'

$y := \alpha * A' * x + \beta * y$

*TRANS* = 'C' or 'c'

$y := \alpha * A' * x + \beta * y$

Unchanged on exit.

**M** On entry, *M* specifies the number of rows of the matrix *A*; *M* must be at least 0; unchanged on exit.

**N** On entry, *N* specifies the number of columns of the matrix *A*; *N* must be at least 0; unchanged on exit.

**KL** On entry, *KL* specifies the number of subdiagonals of the matrix *A*; *KL* must satisfy 0 .le. *KL*; unchanged on exit.

**KU** On entry, *KU* specifies the number of superdiagonals of the matrix *A*; *KU* must satisfy 0 .le. *KU*; unchanged on exit.

**ALPHA** On entry, *ALPHA* specifies the scalar alpha; unchanged on exit.

*A* A vector of dimension ( *LDA*, *N* ); on entry, the leading (  $KL + KU + 1$  ) by *N* part of the array *A* must contain the matrix of coefficients, supplied column by column, with the leading diagonal of the matrix in row (  $KU + 1$  ) of the array, the first superdiagonal starting at position 2 in row *KU*, the first subdiagonal starting at position 1 in row (  $KU + 2$  ), and so on. Elements in the array *A* that do not correspond to elements in the band matrix (such as the top left *KU* by *KU* triangle) are not referenced. The following program segment transfers a band matrix from conventional full matrix storage to band storage:

```

DO 20, J = 1, N
  K = KU + 1 - J
  DO 10, I = MAX( 1, J - KU ), MIN( M, J + KL )
    A( K + I, J ) = matrix( I, J )
  10 CONTINUE
  20 CONTINUE

```

Unchanged on exit.

*LDA* On entry, *LDA* specifies the first dimension of *A* as declared in the calling (sub) program. *LDA* must be at least (  $KL + KU + 1$  ); unchanged on exit.

*X* A vector of dimension at least (  $1 + (N-1) * \text{abs}( INCX )$  ) when *TRANS* = 'N' or 'n', otherwise, at least (  $1 + (M-1) * \text{abs}( INCX )$  ); on entry, the incremented array *X* must contain the vector x; unchanged on exit.

*INCX* On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must not be 0; unchanged on exit.

*BETA* On entry, *BETA* specifies the scalar beta; when *BETA* is supplied as 0 then *Y* need not be set on input; unchanged on exit.

*Y* A vector of dimension at least (  $1 + (M-1) * \text{abs}( INCY )$  ) when *TRANS* = 'N' or 'n', otherwise, at least (  $1 + (N-1) * \text{abs}( INCY )$  ); on entry, the incremented array *Y* must contain the vector y; on exit, *Y* is overwritten by the updated vector y.

*INCY* On entry, *INCY* specifies the increment for the elements of *Y*; *INCY* must not be 0; unchanged on exit.

## CHEMV or ZHEMV Subroutine

### Purpose

Performs matrix-vector operations using Hermitian matrices.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

**SUBROUTINE CHEMV**(*UPLO*, *N*, *ALPHA*, *A*, *LDA*,

*X*, *INCX*, *BETA*, *Y*, *INCY*)

**COMPLEX** *ALPHA*, *BETA*

**INTEGER** *INCX*, *INCY*, *LDA*, *N*

**CHARACTER\*1** *UPLO*

**COMPLEX** *A*(*LDA*,\*), *X*(\*), *Y*(\*)

**SUBROUTINE ZHEMV**(*UPLO*, *N*, *ALPHA*, *A*, *LDA*,

*X*, *INCX*, *BETA*, *Y*, *INCY*)

**COMPLEX\*16** *ALPHA*, *BETA*

**INTEGER** *INCX*, *INCY*, *LDA*, *N*

**CHARACTER\*1** *UPLO*

**COMPLEX\*16** *A*(*LDA*,\*), *X*(\*), *Y*(\*)

### Description

The **CHEMV** or **ZHEMV** subroutine performs the matrix-vector operation:

$$y := \alpha * A * x + \beta * y$$

where alpha and beta are scalars, x and y are *N* element vectors and *A* is an *N* by *N* Hermitian matrix.

## Parameters

<i>UPLO</i>	On entry, <i>UPLO</i> specifies whether the upper or lower triangular part of the array <i>A</i> is to be referenced as follows:  <i>UPLO</i> = 'U' or 'u' Only the upper triangular part of <i>A</i> is to be referenced; unchanged on exit.  <i>UPLO</i> = 'L' or 'l' Only the lower triangular part of <i>A</i> is to be referenced; unchanged on exit.
<i>N</i>	On entry, <i>N</i> specifies the order of the matrix <i>A</i> ; <i>N</i> must be at least 0; unchanged on exit.
<i>ALPHA</i>	On entry, <i>ALPHA</i> specifies the scalar alpha; unchanged on exit.
<i>A</i>	An array of dimension ( <i>LDA</i> , <i>N</i> ); on entry with <i>UPLO</i> = 'U' or 'u', the leading <i>N</i> by <i>N</i> upper triangular part of the array <i>A</i> must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of <i>A</i> is not referenced; on entry with <i>UPLO</i> = 'L' or 'l', the leading <i>N</i> by <i>N</i> lower triangular part of the array <i>A</i> must contain the lower triangular part of the Hermitian matrix and the strictly upper triangular part of <i>A</i> is not referenced. The imaginary parts of the diagonal elements need not be set and are assumed to be 0; unchanged on exit.
<i>LDA</i>	On entry, <i>LDA</i> specifies the first dimension of <i>A</i> as declared in the calling (sub) program; <i>LDA</i> must be at least max( 1, <i>N</i> ); unchanged on exit.
<i>X</i>	A vector of dimension at least (1 + ( <i>N</i> -1) * abs( <i>INCX</i> ) ); on entry, the incremented array <i>X</i> must contain the <i>N</i> element vector <i>x</i> ; unchanged on exit.
<i>INCX</i>	On entry, <i>INCX</i> specifies the increment for the elements of <i>X</i> ; <i>INCX</i> must not be 0; unchanged on exit.
<i>BETA</i>	On entry, <i>BETA</i> specifies the scalar beta; when <i>BETA</i> is supplied as 0 then <i>Y</i> need not be set on input; unchanged on exit.
<i>Y</i>	A vector of dimension at least (1 + ( <i>N</i> -1) * abs( <i>INCX</i> ) ); on entry, the incremented array <i>Y</i> must contain the <i>N</i> element vector <i>y</i> ; on exit, <i>Y</i> is overwritten by the updated vector <i>y</i> .
<i>INCY</i>	On entry, <i>INCY</i> specifies the increment for the elements of <i>Y</i> ; <i>INCY</i> must not be 0; unchanged on exit.

---

## CHBMV or ZHBMV Subroutine

### Purpose

Performs matrix-vector operations using a Hermitian band matrix.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE CHBMV(UPLO, N, K, ALPHA, A, LDA,  
X, INCX, BETA, Y, INCY)
```

```
COMPLEX ALPHA, BETA
```

```
INTEGER INCX, INCY, K, LDA, N
```

```
CHARACTER*1 UPLO
```

```
COMPLEX A(LDA,*), X(*), Y(*)
```

```
SUBROUTINE ZHBMV(UPLO, N, K, ALPHA, A, LDA,  
X, INCX, BETA, Y, INCY)
```

```
COMPLEX*16 ALPHA, BETA
```

```
INTEGER INCX, INCY, K, LDA, N
```

```
CHARACTER*1 UPLO
```

```
COMPLEX*16 A(LDA,*), X(*), Y(*)
```

### Description

The **CHBMV** or **ZHBMV** subroutine performs the matrix-vector operation:

$$y := \alpha * A * x + \beta * y$$

where alpha and beta are scalars, x and y are  $N$  element vectors, and  $A$  is an  $N$  by  $N$  Hermitian band matrix with  $K$  superdiagonals.

## Parameters

*UPLO* On entry, *UPLO* specifies whether the upper or lower triangular part of the band matrix  $A$  is being supplied as follows:

*UPLO* = 'U' or 'u'

The upper triangular part of  $A$  is being supplied.

*UPLO* = 'L' or 'l'

The lower triangular part of  $A$  is being supplied.

Unchanged on exit.

$N$  On entry,  $N$  specifies the order of the matrix  $A$ ;  $N$  must be at least 0; unchanged on exit.

$K$  On entry,  $K$  specifies the number of superdiagonals of the matrix  $A$ ;  $K$  must satisfy  $0 \leq K$ ; unchanged on exit.

*ALPHA* On entry, *ALPHA* specifies the scalar alpha; unchanged on exit.

$A$  An array of dimension (  $LDA$ ,  $N$  ). On entry with *UPLO* = 'U' or 'u', the leading (  $K + 1$  ) by  $N$  part of the array  $A$  must contain the upper triangular band part of the Hermitian matrix, supplied column by column, with the leading diagonal of the matrix in row (  $K + 1$  ) of the array, the first superdiagonal starting at position 2 in row  $K$ , and so on. The top left  $K$  by  $K$  triangle of the array  $A$  is not referenced. The following program segment transfers the upper triangular part of a Hermitian band matrix from conventional full matrix storage to band storage:

```
DO 20, J = 1, N
  M = K + 1 - J
  DO 10, I = MAX( 1, J - K ), J
    A( M + I, J ) = matrix( I, J )
  10 CONTINUE
  20 CONTINUE
```

**Note:** On entry with *UPLO* = 'L' or 'l', the leading (  $K + 1$  ) by  $N$  part of the array  $A$  must contain the lower triangular band part of the Hermitian matrix, supplied column by column, with the leading diagonal of the matrix in row 1 of the array, the first subdiagonal starting at position 1 in row 2, and so on. The bottom right  $K$  by  $K$  triangle of the array  $A$  is not referenced. The following program segment transfers the lower triangular part of a Hermitian band matrix from conventional full matrix storage to band storage:

```
DO 20, J = 1, N
  M = 1 - J
  DO 10, I = J, MIN( N, J + K )
    A( M + I, J ) = matrix( I, J )
  10 CONTINUE
  20 CONTINUE
```

The imaginary parts of the diagonal elements need not be set and are assumed to be 0. Unchanged on exit.

*LDA* On entry, *LDA* specifies the first dimension of  $A$  as declared in the calling (sub) program; *LDA* must be at least (  $K + 1$  ); unchanged on exit.

$X$  A vector of dimension at least (  $1 + (N-1) * \text{abs}( INCX )$  ); on entry, the incremented array  $X$  must contain the vector x; unchanged on exit.

*INCX* On entry, *INCX* specifies the increment for the elements of  $X$ ; *INCX* must not be 0 unchanged on exit.

*BETA* On entry, *BETA* specifies the scalar beta unchanged on exit.

$Y$  A vector of dimension at least (  $1 + (N-1) * \text{abs}( INCY )$  ); on entry, the incremented array  $Y$  must contain the vector y; on exit,  $Y$  is overwritten by the updated vector y.

*INCY* On entry, *INCY* specifies the increment for the elements of  $Y$ ; *INCY* must not be 0; unchanged on exit.

---

## CHPMV or ZHPMV Subroutine

### Purpose

Performs matrix-vector operations using a packed Hermitian matrix.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE CHPMV(UPLO, N, ALPHA, AP, X,  
INCX, BETA, Y, INCY)
```

```
COMPLEX ALPHA, BETA
```

```
INTEGER INCX, INCY, N
```

```
CHARACTER*1 UPLO
```

```
COMPLEX AP(*), X(*), Y(*)
```

```
SUBROUTINE ZHPMV
```

```
COMPLEX*16 ALPHA, BETA
```

```
INTEGER INCX, INCY, N
```

```
CHARACTER*1 UPLO
```

```
COMPLEX*16 AP(*), X(*), Y(*)
```

### Description

The **CHPMV** or **ZHPMV** subroutine performs the matrix-vector operation:

$$y := \alpha * A * x + \beta * y$$

where  $\alpha$  and  $\beta$  are scalars,  $x$  and  $y$  are  $N$  element vectors and  $A$  is an  $N$  by  $N$  Hermitian matrix, supplied in packed form.

### Parameters

**UPLO** On entry, **UPLO** specifies whether the upper or lower triangular part of the matrix  $A$  is supplied in the packed array **AP** as follows:

**UPLO** = 'U' or 'u'

The upper triangular part of  $A$  is supplied in **AP**.

**UPLO** = 'L' or 'l'

The lower triangular part of  $A$  is supplied in **AP**.

Unchanged on exit.

**N** On entry, **N** specifies the order of the matrix  $A$ ; **N** must be at least 0; unchanged on exit.

**ALPHA** On entry, **ALPHA** specifies the scalar  $\alpha$ ; unchanged on exit.

**AP** A vector of dimension at least  $((N * (N+1)) / 2)$ ; on entry with **UPLO** = 'U' or 'u', the array **AP** must contain the upper triangular part of the Hermitian matrix packed sequentially, column by column, so that **AP**(1) contains  $A(1,1)$ , **AP**(2) and **AP**(3) contain  $A(1,2)$  and  $A(2,2)$  respectively, and so on; on entry with **UPLO** = 'L' or 'l', the array **AP** must contain the lower triangular part of the Hermitian matrix packed sequentially, column by column, so that **AP**(1) contains  $A(1,1)$ , **AP**(2) and **AP**(3) contain  $A(2,1)$  and  $A(3,1)$  respectively, and so on. The imaginary parts of the diagonal elements need not be set and are assumed to be 0; unchanged on exit.

**X** A vector of dimension at least  $(1 + (N-1) * \text{abs}( \text{INCX} ))$ ; on entry, the incremented array **X** must contain the  $N$  element vector  $x$ ; unchanged on exit.

**INCX** On entry, **INCX** specifies the increment for the elements of **X**; **INCX** must not be 0; unchanged on exit.

**BETA** On entry, **BETA** specifies the scalar  $\beta$ ; when **BETA** is supplied as 0 then **Y** need not be set on input; unchanged on exit.

*Y* A vector of dimension at least  $(1 + (N-1) * \text{abs}( \text{INCY} ))$ ; on entry, the incremented array *Y* must contain the *N* element vector *y*; on exit, *Y* is overwritten by the updated vector *y*.  
*INCY* On entry, *INCY* specifies the increment for the elements of *Y*; *INCY* must not be 0; unchanged on exit.

---

## SSYMV or DSYMV Subroutine

### Purpose

Performs matrix-vector operations using a symmetric matrix.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE SSYMV(UPLO, N, ALPHA, A, LDA,  
X, INCX, BETA, Y, INCY)
```

```
REAL ALPHA, BETA
```

```
INTEGER INCX, INCY, LDA, N
```

```
CHARACTER*1 UPLO
```

```
REAL A(LDA,*), X(*), Y(*)
```

```
SUBROUTINE DSYMV(UPLO, N, ALPHA, A, LDA,  
X, INCX, BETA, Y, INCY)
```

```
DOUBLE PRECISION ALPHA, BETA
```

```
INTEGER INCX, INCY, LDA, N
```

```
CHARACTER*1 UPLO
```

```
DOUBLE PRECISION A(LDA,*), X(*), Y(*)
```

### Description

The **SSYMV** or **DSYMV** subroutine performs the matrix-vector operation:

$$y := \alpha * A * x + \beta * y$$

where  $\alpha$  and  $\beta$  are scalars,  $x$  and  $y$  are  $N$  element vectors and  $A$  is an  $N$  by  $N$  symmetric matrix.

### Parameters

*UPLO* On entry, *UPLO* specifies whether the upper or lower triangular part of the array *A* is to be referenced as follows:

*UPLO* = 'U' or 'u'

Only the upper triangular part of *A* is to be referenced.

*UPLO* = 'L' or 'l'

Only the lower triangular part of *A* is to be referenced.

Unchanged on exit.

*N* On entry, *N* specifies the order of the matrix *A*; *N* must be at least 0; unchanged on exit.

*ALPHA* On entry, *ALPHA* specifies the scalar  $\alpha$ ; unchanged on exit.

*A* An array of dimension  $(LDA, N)$ ; on entry with *UPLO* = 'U' or 'u', the leading  $N$  by  $N$  upper triangular part of the array *A* must contain the upper triangular part of the symmetric matrix; the strictly lower triangular part of *A* is not referenced; on entry with *UPLO* = 'L' or 'l', the leading  $N$  by  $N$  lower triangular part of the array *A* must contain the lower triangular part of the symmetric matrix; the strictly upper triangular part of *A* is not referenced; unchanged on exit.

*LDA* On entry, *LDA* specifies the first dimension of *A* as declared in the calling (sub) program; *LDA* must be at least  $\max(1, N)$ ; unchanged on exit.

<i>X</i>	A vector of dimension at least $(1 + (N-1) * \text{abs}( INCX ))$ ; on entry, the incremented array <i>X</i> must contain the <i>N</i> element vector <i>x</i> ; unchanged on exit.
<i>INCX</i>	On entry, <i>INCX</i> specifies the increment for the elements of <i>X</i> ; <i>INCX</i> must not be 0; unchanged on exit.
<i>BETA</i>	On entry, <i>BETA</i> specifies the scalar beta; when <i>BETA</i> is supplied as 0 then <i>Y</i> need not be set on input; unchanged on exit.
<i>Y</i>	A vector of dimension at least $(1 + (N-1) * \text{abs}( INCY ))$ ; on entry, the incremented array <i>Y</i> must contain the <i>N</i> element vector <i>y</i> ; on exit, <i>Y</i> is overwritten by the updated vector <i>y</i> .
<i>INCY</i>	On entry, <i>INCY</i> specifies the increment for the elements of <i>Y</i> ; <i>INCY</i> must not be 0; unchanged on exit.

## SSBMV or DSBMV Subroutine

### Purpose

Performs matrix-vector operations using symmetric band matrix.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE SSBMV(UPLO, N, K, ALPHA, A, LDA,  
X, INCX, BETA, Y, INCY)
```

```
REAL ALPHA, BETA
```

```
INTEGER INCX, INCY, K, LDA, N
```

```
CHARACTER*1 UPLO
```

```
REAL A(LDA,*), X(*), Y(*)
```

```
SUBROUTINE DSBMV(UPLO, N, K, ALPHA, A, LDA,  
X, INCX, BETA, Y, INCY)
```

```
DOUBLE PRECISION ALPHA, BETA
```

```
INTEGER INCX, INCY, K, LDA, N
```

```
CHARACTER*1 UPLO
```

```
DOUBLE PRECISION A(LDA,*), X(*), Y(*)
```

### Description

The **SSBMV** or **DSBMV** subroutine performs the matrix-vector operation:

$$y := \alpha * A * x + \beta * y$$

where  $\alpha$  and  $\beta$  are scalars, *x* and *y* are *N* element vectors, and *A* is an *N* by *N* symmetric band matrix with *K* super-diagonals.

### Parameters

*UPLO* On entry, *UPLO* specifies whether the upper or lower triangular part of the band matrix *A* is being supplied as follows:

*UPLO* = 'U' or 'u'

The upper triangular part of *A* is being supplied.

*UPLO* = 'L' or 'l'

The lower triangular part of *A* is being supplied.

Unchanged on exit.

*N* On entry, *N* specifies the order of the matrix *A*; *N* must be at least 0; unchanged on exit.

*K* On entry, *K* specifies the number of superdiagonals of the matrix *A*; *K* must satisfy  $0 \leq K$ ; unchanged on exit.

*ALPHA* On entry, *ALPHA* specifies the scalar  $\alpha$ ; unchanged on exit.

*A* An array of dimension ( *LDA*, *N* ); on entry with *UPLO* = 'U' or 'u', the leading ( *K* + 1 ) by *N* part of the array *A* must contain the upper triangular band part of the symmetric matrix, supplied column by column, with the leading diagonal of the matrix in row ( *K* + 1 ) of the array, the first superdiagonal starting at position 2 in row *K*, and so on. The top left *K* by *K* triangle of the array *A* is not referenced. The following program segment transfers the upper triangular part of a symmetric band matrix from conventional full matrix storage to band storage:

```

DO 20, J = 1, N
    M = K + 1 - J
    DO 10, I = MAX( 1, J - K ), J
        A( M + I, J ) = matrix( I, J )
10 CONTINUE
20 CONTINUE

```

On entry with *UPLO* = 'L' or 'l', the leading ( *K* + 1 ) by *N* part of the array *A* must contain the lower triangular band part of the symmetric matrix, supplied column by column, with the leading diagonal of the matrix in row 1 of the array, the first subdiagonal starting at position 1 in row 2, and so on. The bottom right *K* by *K* triangle of the array *A* is not referenced. The following program segment transfers the lower triangular part of a symmetric band matrix from conventional full matrix storage to band storage:

```

DO 20, J = 1, N
    M = 1 - J
    DO 10, I = J, MIN( N, J + K )
        A( M + I, J ) = matrix( I, J )
10 CONTINUE
20 CONTINUE

```

Unchanged on exit.

*LDA* On entry, *LDA* specifies the first dimension of *A* as declared in the calling (sub) program; *LDA* must be at least ( *K* + 1 ); unchanged on exit.

*X* A vector of dimension at least ( 1 + (*N*-1) \* abs( *INCX* ) ); on entry, the incremented array *X* must contain the vector *x*; unchanged on exit.

*INCX* On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must not be 0; unchanged on exit.

*BETA* On entry, *BETA* specifies the scalar beta; unchanged on exit.

*Y* A vector of dimension at least ( 1 + (*N*-1) \* abs( *INCY* ) ); on entry, the incremented array *Y* must contain the vector *y*; on exit, *Y* is overwritten by the updated vector *y*.

*INCY* On entry, *INCY* specifies the increment for the elements of *Y*; *INCY* must not be 0; unchanged on exit.

## SSPMV or DSPMV Subroutine

### Purpose

Performs matrix-vector operations using a packed symmetric matrix.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```

SUBROUTINE SSPMV(UPLO, N, ALPHA, AP, X,
INCX, BETA, Y, INCY)
REAL ALPHA, BETA
INTEGER INCX, INCY, N
CHARACTER*1 UPLO
REAL AP(*), X(*), Y(*)

```

```

SUBROUTINE DSPMV(UPLO, N, ALPHA, AP, X,
INCX, BETA, Y, INCY)
DOUBLE PRECISION ALPHA,BETA
INTEGER INCX,INCY,N
CHARACTER*1 UPLO
DOUBLE PRECISION AP(*), X(*), Y(*)

```

## Description

The **SSPMV** or **DSPMV** subroutine performs the matrix-vector operation:

$$y := \alpha * A * x + \beta * y$$

where  $\alpha$  and  $\beta$  are scalars,  $x$  and  $y$  are  $N$  element vectors and  $A$  is an  $N$  by  $N$  symmetric matrix, supplied in packed form.

## Parameters

- UPLO** On entry, *UPLO* specifies whether the upper or lower triangular part of the matrix *A* is supplied in the packed array *AP* as follows:
- UPLO = 'U' or 'u'**  
The upper triangular part of *A* is supplied in *AP*.
- UPLO = 'L' or 'l'**  
The lower triangular part of *A* is supplied in *AP*.
- N** Unchanged on exit.
- N** On entry, *N* specifies the order of the matrix *A*; *N* must be at least 0; unchanged on exit.
- ALPHA** On entry, *ALPHA* specifies the scalar  $\alpha$ ; unchanged on exit.
- AP** A vector of dimension at least  $( ( N * (N+1) ) / 2 )$ ; on entry with *UPLO* = 'U' or 'u', the array *AP* must contain the upper triangular part of the symmetric matrix packed sequentially, column by column, so that *AP*(1) contains *A*(1,1), *AP*(2) and *AP*(3) contain *A*(1,2) and *A*(2,2) respectively, and so on; on entry with *UPLO* = 'L' or 'l', the array *AP* must contain the lower triangular part of the symmetric matrix packed sequentially, column by column, so that *AP*(1) contains *A*(1,1), *AP*(2) and *AP*(3) contain *A*(2,1) and *A*(3,1) respectively, and so on; unchanged on exit.
- X** A vector of dimension at least  $( 1 + (N-1) * \text{abs}( INCX ) )$ ; on entry, the incremented array *X* must contain the *N* element vector *x*; unchanged on exit.
- INCX** On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must not be 0; unchanged on exit.
- BETA** On entry, *BETA* specifies the scalar  $\beta$ ; when *BETA* is supplied as 0 then *Y* need not be set on input; unchanged on exit.
- Y** A vector of dimension at least  $( 1 + (N-1) * \text{abs}( INCY ) )$ ; on entry, the incremented array *Y* must contain the *N* element vector *y*; on exit, *Y* is overwritten by the updated vector *y*.
- INCY** On entry, *INCY* specifies the increment for the elements of *Y*; *INCY* must not be 0; unchanged on exit.

---

## STRMV, DTRMV, CTRMV, or ZTRMV Subroutine

### Purpose

Performs matrix-vector operations using a triangular matrix.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```

SUBROUTINE STRMV(UPLO, TRANS, DIAG, N,
A, LDA, X, INCX)

```

```

INTEGER INCX, LDA, N
CHARACTER*1 DIAG, TRANS, UPLO
REAL A(LDA,*), X(*)
SUBROUTINE DTRMV(UPLO, TRANS, DIAG, N,
A, LDA, X, INCX)
INTEGER INCX,LDA,N
CHARACTER*1 DIAG,TRANS,UPLO
DOUBLE PRECISION A(LDA,*), X(*)
SUBROUTINE CTRMV(UPLO, TRANS, DIAG, N,
A, LDA, X, INCX)
INTEGER INCX,LDA,N
CHARACTER*1 DIAG,TRANS,UPLO
COMPLEX A(LDA,*), X(*)
SUBROUTINE ZTRMV(UPLO, TRANS, DIAG, N,
A, LDA, X, INCX)
INTEGER INCX,LDA,N
CHARACTER*1 DIAG,TRANS,UPLO
COMPLEX*16 A(LDA,*),X(*)

```

## Description

The **STRMV**, **DTRMV**, **CTRMV**, or **ZTRMV** subroutine performs one of the matrix-vector operations:

$x := A * x$

OR

$x := A' * x$

where  $x$  is an  $N$  element vector and  $A$  is an  $N$  by  $N$  unit, or non-unit, upper or lower triangular matrix.

## Parameters

**UPLO** On entry, **UPLO** specifies whether the matrix is an upper or lower triangular matrix as follows:

**UPLO = 'U' or 'u'**

$A$  is an upper triangular matrix.

**UPLO = 'L' or 'l'**

$A$  is a lower triangular matrix.

Unchanged on exit.

**TRANS** On entry, **TRANS** specifies the operation to be performed as follows:

**TRANS = 'N' or 'n'**

$x := A * x$

**TRANS = 'T' or 't'**

$x := A' * x$

**TRANS = 'C' or 'c'**

$x := A' * x$

Unchanged on exit.

**DIAG** On entry, **DIAG** specifies whether or not  $A$  is unit triangular as follows:

**DIAG = 'U' or 'u'**

$A$  is assumed to be unit triangular.

**DIAG = 'N' or 'n'**

$A$  is not assumed to be unit triangular.

Unchanged on exit.

<i>N</i>	On entry, <i>N</i> specifies the order of the matrix <i>A</i> ; <i>N</i> must be at least 0; unchanged on exit.
<i>A</i>	An array of dimension ( <i>LDA</i> , <i>N</i> ); on entry with <i>UPLO</i> = 'U' or 'u', the leading <i>N</i> by <i>N</i> upper triangular part of the array <i>A</i> must contain the upper triangular matrix and the strictly lower triangular part of <i>A</i> is not referenced; on entry with <i>UPLO</i> = 'L' or 'l', the leading <i>N</i> by <i>N</i> lower triangular part of the array <i>A</i> must contain the lower triangular matrix and the strictly upper triangular part of <i>A</i> is not referenced. When <i>DIAG</i> = 'U' or 'u', the diagonal elements of <i>A</i> are not referenced, but are assumed to be unity; unchanged on exit.
<i>LDA</i>	On entry, <i>LDA</i> specifies the first dimension of <i>A</i> as declared in the calling (sub) program. <i>LDA</i> must be at least max( 1, <i>N</i> ); unchanged on exit.
<i>X</i>	A vector of dimension at least (1 + ( <i>N</i> -1) * abs( <i>INCX</i> ) ). On entry, the incremented array <i>X</i> must contain the <i>N</i> element vector <i>x</i> ; on exit, <i>X</i> is overwritten with the transformed vector <i>x</i> .
<i>INCX</i>	On entry, <i>INCX</i> specifies the increment for the elements of <i>X</i> ; <i>INCX</i> must not be 0; unchanged on exit.

## STBMV, DTBMV, CTBMV, or ZTBMV Subroutine

### Purpose

Performs matrix-vector operations using a triangular band matrix.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE STBMV(UPLO, TRANS, DIAG, N,  
K, A, LDA, X, INCX)
```

```
INTEGER INCX, K, LDA, N
```

```
CHARACTER*1 DIAG, TRANS, UPLO
```

```
REAL A(LDA,*), X(*)
```

```
SUBROUTINE DTBMV(UPLO, TRANS, DIAG, N,  
K, A, LDA, X, INCX)
```

```
INTEGER INCX, K, LDA, N
```

```
CHARACTER*1 DIAG, TRANS, UPLO
```

```
DOUBLE PRECISION A(LDA,*), X(*)
```

```
SUBROUTINE CTBMV(UPLO, TRANS, DIAG, N,  
K, A, LDA, X, INCX)
```

```
INTEGER INCX, K, LDA, N
```

```
CHARACTER*1 DIAG, TRANS, UPLO
```

```
COMPLEX A(LDA,*), X(*)
```

```
SUBROUTINE ZTBMV(UPLO, TRANS, DIAG, N,  
K, A, LDA, X, INCX)
```

```
INTEGER INCX, K, LDA, N
```

```
CHARACTER*1 DIAG, TRANS, UPLO
```

```
COMPLEX*16 A(LDA,*), X(*)
```

### Description

The **STBMV**, **DTBMV**, **CTBMV**, or **ZTBMV** subroutine performs one of the matrix-vector operations:

```
x := A * x
```

OR

```
x := A' * x
```

where *x* is an *N* element vector and *A* is an *N* by *N* unit, or non-unit, upper or lower triangular band matrix, with ( *K* + 1 ) diagonals.

## Parameters

*UPLO* On entry, *UPLO* specifies whether the matrix is an upper or lower triangular matrix as follows:

*UPLO* = 'U' or 'u'

A is an upper triangular matrix.

*UPLO* = 'L' or 'l'

A is a lower triangular matrix.

Unchanged on exit.

*TRANS* On entry, *TRANS* specifies the operation to be performed as follows:

*TRANS* = 'N' or 'n'

$x := A * x$

*TRANS* = 'T' or 't'

$x := A' * x$

*TRANS* = 'C' or 'c'

$x := A' * x$

Unchanged on exit.

*DIAG* On entry, *DIAG* specifies whether or not *A* is unit triangular as follows:

*DIAG* = 'U' or 'u'

A is assumed to be unit triangular.

*DIAG* = 'N' or 'n'

A is not assumed to be unit triangular.

Unchanged on exit.

*N* On entry, *N* specifies the order of the matrix *A*; *N* must be at least 0; unchanged on exit.

*K* On entry with *UPLO* = 'U' or 'u', *K* specifies the number of superdiagonals of the matrix *A*; on entry with *UPLO* = 'L' or 'l', *K* specifies the number of subdiagonals of the matrix *A*. *K* must satisfy  $0 \leq K$ ; unchanged on exit.

*A* An array of dimension ( *LDA*, *N* ). On entry with *UPLO* = 'U' or 'u', the leading ( *K* + 1 ) by *N* part of the array *A* must contain the upper triangular band part of the matrix of coefficients, supplied column by column, with the leading diagonal of the matrix in row ( *K* + 1 ) of the array, the first superdiagonal starting at position 2 in row *K*, and so on. The top left *K* by *K* triangle of the array *A* is not referenced. The following program segment will transfer an upper triangular band matrix from conventional full matrix storage to band storage:

```
DO 20, J = 1, N
    M = K + 1 - J
    DO 10, I = MAX( 1, J - K ), J
        A( M + I, J ) = matrix( I, J )
    10 CONTINUE
20 CONTINUE
```

```
DO 20, J = 1, N
    M = 1 - J
    DO 10, I = J, MIN( N, J + K )
        A( M + I, J ) = matrix( I, J )
    10 CONTINUE
20 CONTINUE
```

On entry with *UPLO* = 'L' or 'l', the leading ( *K* + 1 ) by *N* part of the array *A* must contain the lower triangular band part of the matrix of coefficients, supplied column by column, with the leading diagonal of the matrix in row 1 of the array, the first subdiagonal starting at position 1 in row 2, and so on. The bottom right *K* by *K* triangle of the array *A* is not referenced. The following program segment will transfer a lower triangular band matrix from conventional full matrix storage to band storage:

When *DIAG* = 'U' or 'u' the elements of the array *A* corresponding to the diagonal elements of the matrix are not referenced, but are assumed to be unity; unchanged on exit.

*LDA* On entry, *LDA* specifies the first dimension of *A* as declared in the calling (sub) program; *LDA* must be at least ( *K* + 1 ); unchanged on exit.

*X* A vector of dimension at least (1 + (*N*-1) \* abs( *INCX* ) ); on entry, the incremented array *X* must contain the *N* element vector *x*; on exit, *X* is overwritten with the transformed vector *x*.

*INCX* On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must not be 0; unchanged on exit.

## STPMV, DTPMV, CTPMV, or ZTPMV Subroutine

### Purpose

Performs matrix-vector operations on a packed triangular matrix.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE STPMV(UPLO, TRANS, DIAG,
N, AP, X, INCX)
```

```
INTEGER INCX, N
```

```
CHARACTER*1 DIAG, TRANS, UPLO
```

```
REAL AP(*), X(*)
```

```
SUBROUTINE DTPMV(UPLO, TRANS, DIAG,
N, AP, X, INCX)
```

```
INTEGER INCX, N
```

```
CHARACTER*1 DIAG, TRANS, UPLO
```

```
DOUBLE PRECISION AP(*), X(*)
```

```
SUBROUTINE CTPMV(UPLO, TRANS, DIAG,
N, AP, X, INCX)
```

```
INTEGER INCX, N
```

```
CHARACTER*1 DIAG, TRANS, UPLO
```

```
COMPLEX AP(*), X(*)
```

```
SUBROUTINE ZTPMV(UPLO, TRANS, DIAG,
N, AP, X, INCX)
```

```
INTEGER INCX, N
```

```
CHARACTER*1 DIAG, TRANS, UPLO
```

```
COMPLEX*16 AP(*), X(*)
```

### Description

The **STPMV**, **DTPMV**, **CTPMV**, or **ZTPMV** subroutine performs one of the matrix-vector operations:

```
x := A * x
```

OR

```
x := A' * x
```

where *x* is an *N* element vector and *A* is an *N* by *N* unit, or non-unit, upper or lower triangular matrix, supplied in packed form.

## Parameters

- UPLO* On entry, *UPLO* specifies whether the matrix is an upper or lower triangular matrix as follows:
- UPLO* = 'U' or 'u'  
A is an upper triangular matrix.
- UPLO* = 'L' or 'l'  
A is a lower triangular matrix.
- Unchanged on exit.
- TRANS* On entry, *TRANS* specifies the operation to be performed as follows:
- TRANS* = 'N' or 'n'  
 $x := A * x$
- TRANS* = 'T' or 't'  
 $x := A' * x$
- TRANS* = 'C' or 'c'  
 $x := A' * x$
- Unchanged on exit.
- DIAG* On entry, *DIAG* specifies whether or not *A* is unit triangular as follows:
- DIAG* = 'U' or 'u'  
A is assumed to be unit triangular.
- DIAG* = 'N' or 'n'  
A is not assumed to be unit triangular.
- Unchanged on exit.
- N* On entry, *N* specifies the order of the matrix *A*; *N* must be at least 0; unchanged on exit.
- AP* A vector of dimension at least  $( ( N * (N+1) ) / 2 )$ . On entry with *UPLO* = 'U' or 'u', the array *AP* must contain the upper triangular matrix packed sequentially, column by column, so that *AP*(1) contains *A*(1,1), *AP*(2) and *AP*(3) contain *A*(1,2) and *A*(2,2) respectively, and so on. On entry with *UPLO* = 'L' or 'l', the array *AP* must contain the lower triangular matrix packed sequentially, column by column, so that *AP*(1) contains *A*(1,1), *AP*(2) and *AP*(3) contain *A*(2,1) and *A*(3,1) respectively, and so on. When *DIAG* = 'U' or 'u', the diagonal elements of *A* are not referenced, but are assumed to be unity; unchanged on exit.
- X* A vector of dimension at least  $( 1 + (N-1) * \text{abs}( INCX ) )$ ; on entry, the incremented array *X* must contain the *N* element vector *x*; on exit, *X* is overwritten with the transformed vector *x*.
- INCX* On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must not be 0; unchanged on exit.

---

## STRSV, DTRSV, CTRSV, or ZTRSV Subroutine

### Purpose

Solves system of equations.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE STRSV(UPLO, TRANS, DIAG,  
N, A, LDA, X, INCX)  
INTEGER INCX, LDA, N  
CHARACTER*1 DIAG, TRANS, UPLO  
REAL A(LDA,*), X(*)
```

```

SUBROUTINE DTRSV(UPLO, TRANS, DIAG,
N, A, LDA, X, INCX)
INTEGER INCX,LDA,N
CHARACTER*1 DIAG,TRANS,UPLO
DOUBLE PRECISION A(LDA,*), X(*)
SUBROUTINE CTRSV(UPLO, TRANS, DIAG,
N, A, LDA, X, INCX)
INTEGER INCX,LDA,N
CHARACTER*1 DIAG,TRANS,UPLO
COMPLEX A(LDA,*), X(*)
SUBROUTINE ZTRSV(UPLO, TRANS, DIAG,
N, A, LDA, X, INCX)
INTEGER INCX,LDA,N
CHARACTER*1 DIAG,TRANS,UPLO
COMPLEX*16 A(LDA,*), X(*)

```

## Description

The **STRSV**, **DTRSV**, **CTRSV**, or **ZTRSV** subroutine solves one of the systems of equations:

$$A * x = b$$

OR

$$A' * x = b$$

where  $b$  and  $x$  are  $N$  element vectors and  $A$  is an  $N$  by  $N$  unit, or non-unit, upper or lower triangular matrix.

No test for singularity or near-singularity is included in this routine. Such tests must be performed before calling this routine.

## Parameters

**UPLO** On entry, **UPLO** specifies whether the matrix is an upper or lower triangular matrix as follows:

**UPLO = 'U' or 'u'**  
 $A$  is an upper triangular matrix.

**UPLO = 'L' or 'l'**  
 $A$  is a lower triangular matrix.

Unchanged on exit.

**TRANS** On entry, **TRANS** specifies the equations to be solved as follows:

**TRANS = 'N' or 'n'**  
 $A * x = b$

**TRANS = 'T' or 't'**  
 $A' * x = b$

**TRANS = 'C' or 'c'**  
 $A' * x = b$

Unchanged on exit.

*DIAG* On entry, *DIAG* specifies whether or not *A* is unit triangular as follows:

*DIAG* = 'U' or 'u'  
*A* is assumed to be unit triangular.

*DIAG* = 'N' or 'n'  
*A* is not assumed to be unit triangular.

Unchanged on exit.

*N* On entry, *N* specifies the order of the matrix *A*; *N* must be at least 0; unchanged on exit.

*A* An array of dimension ( *LDA*, *N* ); on entry with *UPLO* = 'U' or 'u', the leading *N* by *N* upper triangular part of the array *A* must contain the upper triangular matrix and the strictly lower triangular part of *A* is not referenced. On entry with *UPLO* = 'L' or 'l', the leading *N* by *N* lower triangular part of the array *A* must contain the lower triangular matrix and the strictly upper triangular part of *A* is not referenced. When *DIAG* = 'U' or 'u', the diagonal elements of *A* are not referenced, but are assumed to be unity; unchanged on exit.

*LDA* On entry, *LDA* specifies the first dimension of *A* as declared in the calling (sub) program; *LDA* must be at least max( 1, *N* ); unchanged on exit.

*X* A vector of dimension at least (1 + (*N*-1) \* abs(*INCX*)); on entry, the incremented array *X* must contain the *N* element right-hand side vector *b*; on exit, *X* is overwritten with the solution vector *x*.

*INCX* On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must not be 0; unchanged on exit.

---

## STBSV, DTBSV, CTBSV, or ZTBSV Subroutine

### Purpose

Solves system of equations.

### Library

BLAS Library (*libblas.a*)

### FORTRAN Syntax

**SUBROUTINE STBSV**(*UPLO*, *TRANS*, *DIAG*,  
*N*, *K*, *A*, *LDA*, *X*, *INCX*)

**INTEGER** *INCX*, *K*, *LDA*, *N*

**CHARACTER\*1** *DIAG*, *TRANS*, *UPLO*

**REAL** *A*(*LDA*,\*), *X*(\*)

**SUBROUTINE DTBSV**(*UPLO*, *TRANS*, *DIAG*,  
*N*, *K*, *A*, *LDA*, *X*, *INCX*)

**INTEGER** *INCX*, *K*, *LDA*, *N*

**CHARACTER\*1** *DIAG*, *TRANS*, *UPLO*

**DOUBLE PRECISION** *A*(*LDA*,\*), *X*(\*)

**SUBROUTINE CTBSV**(*UPLO*, *TRANS*, *DIAG*,  
*N*, *K*, *A*, *LDA*, *X*, *INCX*)

**INTEGER** *INCX*, *K*, *LDA*, *N*

**CHARACTER\*1** *DIAG*, *TRANS*, *UPLO*

**COMPLEX** *A*(*LDA*,\*), *X*(\*)

**SUBROUTINE ZTBSV**(*UPLO*, *TRANS*, *DIAG*,  
*N*, *K*, *A*, *LDA*, *X*, *INCX*)

**INTEGER** *INCX*, *K*, *LDA*, *N*

**CHARACTER\*1** *DIAG*, *TRANS*, *UPLO*

**COMPLEX\*16** *A*(*LDA*,\*), *X*(\*)

### Description

The **STBSV**, **DTBSV**, **CTBSV**, or **ZTBSV** subroutine solves one of the systems of equations:

$$A * x = b$$

OR

$$A' * x = b$$

where  $b$  and  $x$  are  $N$  element vectors and  $A$  is an  $N$  by  $N$  unit, or non-unit, upper or lower triangular band matrix, with  $(K + 1)$  diagonals.

No test for singularity or near-singularity is included in this routine. Such tests must be performed before calling this routine.

## Parameters

*UPLO* On entry, *UPLO* specifies whether the matrix is an upper or lower triangular matrix as follows:

*UPLO* = 'U' or 'u'

$A$  is an upper triangular matrix.

*UPLO* = 'L' or 'l'

$A$  is a lower triangular matrix.

Unchanged on exit.

*TRANS* On entry, *TRANS* specifies the equations to be solved as follows:

*TRANS* = 'N' or 'n'

$A * x = b$

*TRANS* = 'T' or 't'

$A' * x = b$

*TRANS* = 'C' or 'c'

$A' * x = b$

Unchanged on exit.

*DIAG* On entry, *DIAG* specifies whether  $A$  is unit triangular as follows:

*DIAG* = 'U' or 'u'

$A$  is assumed to be unit triangular.

*DIAG* = 'N' or 'n'

$A$  is not assumed to be unit triangular.

Unchanged on exit.

*N* On entry, *N* specifies the order of the matrix  $A$ ; *N* must be at least 0; unchanged on exit.

*K* On entry with *UPLO* = 'U' or 'u', *K* specifies the number of superdiagonals of the matrix  $A$ . On entry with *UPLO* = 'L' or 'l', *K* specifies the number of subdiagonals of the matrix  $A$ ; *K* must satisfy  $0 \leq K$ ; unchanged on exit.

**A** An array of dimension ( *LDA*, *N* ). On entry with *UPLO* = 'U' or 'u', the leading ( *K* + 1 ) by *N* part of the array *A* must contain the upper triangular band part of the matrix of coefficients, supplied column by column, with the leading diagonal of the matrix in row ( *K* + 1 ) of the array, the first superdiagonal starting at position 2 in row *K*, and so on. The top left *K* by *K* triangle of the array *A* is not referenced.

The following program segment will transfer an upper triangular band matrix from conventional full matrix storage to band storage:

```

DO 20, J = 1, N
    M = K + 1 - J
    DO 10, I = MAX( 1, J - K ), J
        A( M + I, J ) = matrix( I, J )
10 CONTINUE
20 CONTINUE

```

On entry with *UPLO* = 'L' or 'l', the leading ( *K* + 1 ) by *N* part of the array *A* must contain the lower triangular band part of the matrix of coefficients, supplied column by column, with the leading diagonal of the matrix in row 1 of the array, the first subdiagonal starting at position 1 in row 2, and so on. The bottom right *K* by *K* triangle of the array *A* is not referenced.

The following program segment will transfer a lower triangular band matrix from conventional full matrix storage to band storage:

```

DO 20, J = 1, N
    M = 1 - J
    DO 10, I = J, MIN( N, J + K )
        A( M + I, J ) = matrix( I, J )
10 CONTINUE
20 CONTINUE

```

When *DIAG* = 'U' or 'u' the elements of the array *A* corresponding to the diagonal elements of the matrix are not referenced, but are assumed to be unity. Unchanged on exit.

**LDA** On entry, *LDA* specifies the first dimension of *A* as declared in the calling (sub) program; *LDA* must be at least ( *K* + 1 ); unchanged on exit.

**X** A vector of dimension at least ( 1 + (*N*-1) \* abs(*INCX*) ); on entry, the incremented array *X* must contain the *N* element right-hand side vector *b*; on exit, *X* is overwritten with the solution vector *x*.

**INCX** On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must not be 0; unchanged on exit.

## STPSV, DTPSV, CTPSV, or ZTPSV Subroutine

### Purpose

Solves systems of equations.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

**SUBROUTINE STPSV**(*UPLO*, *TRANS*, *DIAG*,  
*N*, *AP*, *X*, *INCX*)

**INTEGER** *INCX*, *N*

**CHARACTER\*1** *DIAG*, *TRANS*, *UPLO*

**REAL** *AP*(\*), *X*(\*)

**SUBROUTINE DTPSV**(*UPLO*, *TRANS*, *DIAG*,  
*N*, *AP*, *X*, *INCX*)

**INTEGER** *INCX*, *N*

**CHARACTER\*1** *DIAG*, *TRANS*, *UPLO*

**DOUBLE PRECISION** *AP*(\*), *X*(\*)

```

SUBROUTINE CTPSV(UPLO, TRANS, DIAG,
N, AP, X, INCX)
INTEGER INCX,N
CHARACTER*1 DIAG,TRANS,UPLO
COMPLEX AP(*), X(*)
SUBROUTINE ZTPSV(UPLO, TRANS, DIAG,
N, AP, X, INCX)
INTEGER INCX,N
CHARACTER*1 DIAG,TRANS,UPLO
COMPLEX*16 AP(*), X(*)

```

## Description

The **STPSV**, **DTPSV**, **DTPSV**, or **ZTPSV** subroutine solves one of the systems of equations:

$$A * x = b$$

OR

$$A' * x = b$$

where  $b$  and  $x$  are  $N$  element vectors and  $A$  is an  $N$  by  $N$  unit, or non-unit, upper or lower triangular matrix, supplied in packed form.

No test for singularity or near-singularity is included in this routine. Such tests must be performed before calling this routine.

## Parameters

**UPLO** On entry, *UPLO* specifies whether the matrix is an upper or lower triangular matrix as follows:

*UPLO* = 'U' or 'u'  
 $A$  is an upper triangular matrix.

*UPLO* = 'L' or 'l'  
 $A$  is a lower triangular matrix.

Unchanged on exit.

**TRANS** On entry, *TRANS* specifies the equations to be solved as follows:

*TRANS* = 'N' or 'n'  
 $A * x = b$

*TRANS* = 'T' or 't'  
 $A' * x = b$

*TRANS* = 'C' or 'c'  
 $A' * x = b$

Unchanged on exit.

**DIAG** On entry, *DIAG* specifies whether or not  $A$  is unit triangular as follows:

*DIAG* = 'U' or 'u'  
 $A$  is assumed to be unit triangular.

*DIAG* = 'N' or 'n'  
 $A$  is not assumed to be unit triangular.

Unchanged on exit.

**N** On entry,  $N$  specifies the order of the matrix  $A$ ;  $N$  must be at least 0; unchanged on exit.

<i>AP</i>	A vector of dimension at least $( ( N * (N+1) )/2 )$ ; on entry with <i>UPLO</i> = 'U' or 'u', the array <i>AP</i> must contain the upper triangular matrix packed sequentially, column by column, so that <i>AP</i> (1) contains <i>A</i> (1,1), <i>AP</i> (2) and <i>AP</i> (3) contain <i>A</i> (1,2) and <i>A</i> (2,2) respectively, and so on. Before entry with <i>UPLO</i> = 'L' or 'l', the array <i>AP</i> must contain the lower triangular matrix packed sequentially, column by column, so that <i>AP</i> (1) contains <i>A</i> (1,1), <i>AP</i> (2) and <i>AP</i> (3) contain <i>A</i> (2,1) and <i>A</i> (3,1) respectively, and so on. When <i>DIAG</i> = 'U' or 'u', the diagonal elements of <i>A</i> are not referenced, but are assumed to be unity; unchanged on exit.
<i>X</i>	A vector of dimension at least $( 1 + (N-1) * \text{abs}(INCX) )$ ; on entry, the incremented array <i>X</i> must contain the <i>N</i> element right-hand side vector <i>b</i> ; on exit, <i>X</i> is overwritten with the solution vector <i>x</i> .
<i>INCX</i>	On entry, <i>INCX</i> specifies the increment for the elements of <i>X</i> ; <i>INCX</i> must not be 0; unchanged on exit.

## SGER or DGER Subroutine

### Purpose

Performs the rank 1 operation.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE SGER(M, N, ALPHA, X,  
INCX, Y, INCY, A, LDA)
```

```
REAL ALPHA
```

```
INTEGER INCX, INCY, LDA, M, N
```

```
REAL A(LDA,*), X(*), Y(*)
```

```
SUBROUTINE DGER(M, N, ALPHA, X,  
INCX, Y, INCY, A, LDA)
```

```
DOUBLE PRECISION ALPHA
```

```
INTEGER INCX, INCY, LDA, M, N
```

```
DOUBLE PRECISION A(LDA,*), X(*), Y(*)
```

### Description

The **SGER** or **DGER** subroutine performs the rank 1 operation:

$$A := \alpha * x * y' + A$$

where *alpha* is a scalar, *x* is an *M* element vector, *y* is an *N* element vector and *A* is an *M* by *N* matrix.

### Parameters

<i>M</i>	On entry, <i>M</i> specifies the number of rows of the matrix <i>A</i> ; <i>M</i> must be at least 0; unchanged on exit.
<i>N</i>	On entry, <i>N</i> specifies the number of columns of the matrix <i>A</i> ; <i>N</i> must be at least 0; unchanged on exit.
<i>ALPHA</i>	On entry, <i>ALPHA</i> specifies the scalar <i>alpha</i> ; unchanged on exit.
<i>X</i>	A vector of dimension at least $( 1 + (M-1) * \text{abs}(INCX) )$ ; on entry, the incremented array <i>X</i> must contain the <i>M</i> element vector <i>x</i> ; unchanged on exit.
<i>INCX</i>	On entry, <i>INCX</i> specifies the increment for the elements of <i>X</i> ; <i>INCX</i> must not be 0; unchanged on exit.
<i>Y</i>	A vector of dimension at least $( 1 + (N-1) * \text{abs}(INCY) )$ ; on entry, the incremented array <i>Y</i> must contain the <i>N</i> element vector <i>y</i> ; unchanged on exit.
<i>INCY</i>	On entry, <i>INCY</i> specifies the increment for the elements of <i>Y</i> ; <i>INCY</i> must not be 0; unchanged on exit.
<i>A</i>	An array of dimension $( LDA, N )$ ; on entry, the leading <i>M</i> by <i>N</i> part of the array <i>A</i> must contain the matrix of coefficients; on exit, <i>A</i> is overwritten by the updated matrix.
<i>LDA</i>	On entry, <i>LDA</i> specifies the first dimension of <i>A</i> as declared in the calling (sub) program; <i>LDA</i> must be at least $\max( 1, M )$ ; unchanged on exit.

---

## CGERU or ZGERU Subroutine

### Purpose

Performs the rank 1 operation.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE CGERU(M, N, ALPHA, X, INCX,  
Y, INCY, A, LDA)
```

```
COMPLEX ALPHA
```

```
INTEGER INCX, INCY, LDA, M, N
```

```
COMPLEX A(LDA,*), X(*), Y(*)
```

```
SUBROUTINE ZGERU
```

```
COMPLEX*16 ALPHA
```

```
INTEGER INCX, INCY, LDA, M, N
```

```
COMPLEX*16 A(LDA,*), X(*), Y(*)
```

### Description

The **CGERU** or **ZGERU** subroutine performs the rank 1 operation:

$$A := \alpha * x * y' + A$$

where  $\alpha$  is a scalar,  $x$  is an  $M$  element vector,  $y$  is an  $N$  element vector and  $A$  is an  $M$  by  $N$  matrix.

### Parameters

<i>M</i>	On entry, <i>M</i> specifies the number of rows of the matrix <i>A</i> ; <i>M</i> must be at least 0; unchanged on exit.
<i>N</i>	On entry, <i>N</i> specifies the number of columns of the matrix <i>A</i> ; <i>N</i> must be at least 0; unchanged on exit.
<i>ALPHA</i>	On entry, <i>ALPHA</i> specifies the scalar $\alpha$ ; unchanged on exit.
<i>X</i>	A vector of dimension at least $(1 + (M-1) * \text{abs}(\text{INCX}))$ ; on entry, the incremented array <i>X</i> must contain the $M$ element vector $x$ ; unchanged on exit.
<i>INCX</i>	On entry, <i>INCX</i> specifies the increment for the elements of <i>X</i> ; <i>INCX</i> must not be 0; unchanged on exit.
<i>Y</i>	A vector of dimension at least $(1 + (N-1) * \text{abs}(\text{INCY}))$ ; on entry, the incremented array <i>Y</i> must contain the $N$ element vector $y$ ; unchanged on exit.
<i>INCY</i>	On entry, <i>INCY</i> specifies the increment for the elements of <i>Y</i> ; <i>INCY</i> must not be 0; unchanged on exit.
<i>A</i>	An array of dimension ( <i>LDA</i> , <i>N</i> ); on entry, the leading $M$ by $N$ part of the array <i>A</i> must contain the matrix of coefficients; on exit, <i>A</i> is overwritten by the updated matrix.
<i>LDA</i>	On entry, <i>LDA</i> specifies the first dimension of <i>A</i> as declared in the calling (sub) program; <i>LDA</i> must be at least $\max(1, M)$ ; unchanged on exit.

---

## CGERC or ZGERC Subroutine

### Purpose

Performs the rank 1 operation.

### Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

```
SUBROUTINE CGERC(M, N, ALPHA, X, INCX,  
Y, INCY, A, LDA)  
COMPLEX ALPHA  
INTEGER INCX, INCY, LDA, M, N  
COMPLEX A(LDA,*), X(*), Y(*)  
  
SUBROUTINE ZGERC  
COMPLEX*16 ALPHA  
INTEGER INCX, INCY, LDA, M, N  
COMPLEX*16 A(LDA,*), X(*), Y(*)
```

## Description

The **CGERC** or **ZGERC** subroutine performs the rank 1 operation:

$$A := \alpha * x * \text{conjg}(y^T) + A$$

where  $\alpha$  is a scalar,  $x$  is an  $M$  element vector,  $y$  is an  $N$  element vector and  $A$  is an  $M$  by  $N$  matrix.

## Parameters

<i>M</i>	On entry, <i>M</i> specifies the number of rows of the matrix <i>A</i> ; <i>M</i> must be at least 0; unchanged on exit.
<i>N</i>	On entry, <i>N</i> specifies the number of columns of the matrix <i>A</i> ; <i>N</i> must be at least 0; unchanged on exit.
<i>ALPHA</i>	On entry, <i>ALPHA</i> specifies the scalar $\alpha$ ; unchanged on exit.
<i>X</i>	A vector of dimension at least $(1 + (M-1) * \text{abs}(INCX))$ ; on entry, the incremented array <i>X</i> must contain the $M$ element vector $x$ ; unchanged on exit.
<i>INCX</i>	On entry, <i>INCX</i> specifies the increment for the elements of <i>X</i> ; <i>INCX</i> must not be 0; unchanged on exit.
<i>Y</i>	A vector of dimension at least $(1 + (N-1) * \text{abs}(INCY))$ ; on entry, the incremented array <i>Y</i> must contain the $N$ element vector $y$ ; unchanged on exit.
<i>INCY</i>	On entry, <i>INCY</i> specifies the increment for the elements of <i>Y</i> ; <i>INCY</i> must not be 0; unchanged on exit.
<i>A</i>	An array of dimension $(LDA, N)$ ; on entry, the leading $M$ by $N$ part of the array <i>A</i> must contain the matrix of coefficients; on exit, <i>A</i> is overwritten by the updated matrix.
<i>LDA</i>	On entry, <i>LDA</i> specifies the first dimension of <i>A</i> as declared in the calling (sub) program; <i>LDA</i> must be at least $\max(1, M)$ ; unchanged on exit.

---

## CHER or ZHER Subroutine

### Purpose

Performs the Hermitian rank 1 operation.

### Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

```
SUBROUTINE CHER(UPLO, N, ALPHA,  
X, INCX, A, LDA)  
REAL ALPHA  
INTEGER INCX, LDA, N  
CHARACTER*1 UPLO  
COMPLEX A(LDA,*), X(*)
```

```

SUBROUTINE ZHER(UPLO, N, ALPHA,
X, INCX, A, LDA)
DOUBLE PRECISION ALPHA
INTEGER INCX,LDA,N
CHARACTER*1 UPLO
COMPLEX*16 A(LDA,*), X(*)

```

## Description

The **CHER** or **ZHER** subroutine performs the Hermitian rank 1 operation:

$$A := \alpha * x * \text{conjg}(x^T) + A$$

where  $\alpha$  is a real scalar,  $x$  is an  $N$  element vector and  $A$  is an  $N$  by  $N$  Hermitian matrix.

## Parameters

**UPLO** On entry, *UPLO* specifies whether the upper or lower triangular part of the array *A* is to be referenced as follows:

**UPLO = 'U' or 'u'**  
Only the upper triangular part of *A* is to be referenced.

**UPLO = 'L' or 'l'**  
Only the lower triangular part of *A* is to be referenced.

Unchanged on exit.

**N** On entry, *N* specifies the order of the matrix *A*; *N* must be at least 0; unchanged on exit.

**ALPHA** On entry, *ALPHA* specifies the scalar  $\alpha$ ; unchanged on exit.

**X** A vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; on entry, the incremented array *X* must contain the  $N$  element vector  $x$ ; unchanged on exit.

**INCX** On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must not be 0; unchanged on exit.

**A** An array of dimension  $(LDA, N)$ ; on entry with *UPLO* = 'U' or 'u', the leading  $N$  by  $N$  upper triangular part of the array *A* must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of *A* is not referenced. On exit, the upper triangular part of the array *A* is overwritten by the upper triangular part of the updated matrix. On entry with *UPLO* = 'L' or 'l', the leading  $N$  by  $N$  lower triangular part of the array *A* must contain the lower triangular part of the Hermitian matrix and the strictly upper triangular part of *A* is not referenced. On exit, the lower triangular part of the array *A* is overwritten by the lower triangular part of the updated matrix. The imaginary parts of the diagonal elements need not be set, they are assumed to be 0, and on exit they are set to 0.

**LDA** On entry, *LDA* specifies the first dimension of *A* as declared in the calling (sub) program; *LDA* must be at least  $\max(1, N)$ ; unchanged on exit.

---

## CHPR or ZHPR Subroutine

### Purpose

Performs the Hermitian rank 1 operation.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```

SUBROUTINE CHPR(UPLO, N, ALPHA,
X, INCX, AP)
REAL ALPHA

```

```

INTEGER INCX, N
CHARACTER*1 UPLO
COMPLEX AP(*), X(*)
SUBROUTINE ZHPR(UPLO, N, ALPHA,
X, INCX, AP)
DOUBLE PRECISION ALPHA
INTEGER INCX, N
CHARACTER*1 UPLO
COMPLEX*16 AP(*), X(*)

```

## Description

The **CHPR** or **ZHPR** subroutine performs the Hermitian rank 1 operation:

$$A := \alpha * x * \text{conjg}(x') + A$$

where  $\alpha$  is a real scalar,  $x$  is an  $N$  element vector and  $A$  is an  $N$  by  $N$  Hermitian matrix, supplied in packed form.

## Parameters

**UPLO** On entry, *UPLO* specifies whether the upper or lower triangular part of the matrix  $A$  is supplied in the packed array  $AP$  as follows:

**UPLO = 'U' or 'u'**  
The upper triangular part of  $A$  is supplied in  $AP$ .

**UPLO = 'L' or 'l'**  
The lower triangular part of  $A$  is supplied in  $AP$ .

Unchanged on exit.

**N** On entry,  $N$  specifies the order of the matrix  $A$ ;  $N$  must be at least 0; unchanged on exit.

**ALPHA** On entry, *ALPHA* specifies the scalar  $\alpha$ ; unchanged on exit.

**X** A vector of dimension at least  $(1 + (N-1) * \text{abs}(INCX))$ ; on entry, the incremented array  $X$  must contain the  $N$  element vector  $x$ ; unchanged on exit.

**INCX** On entry, *INCX* specifies the increment for the elements of  $X$ ; *INCX* must not be 0; unchanged on exit.

**AP** A vector of dimension at least  $(N * (N+1) / 2)$ ; on entry with *UPLO* = 'U' or 'u', the array  $AP$  must contain the upper triangular part of the Hermitian matrix packed sequentially, column by column, so that  $AP(1)$  contains  $A(1,1)$ ,  $AP(2)$  and  $AP(3)$  contain  $A(1,2)$  and  $A(2,2)$  respectively, and so on. On exit, the array  $AP$  is overwritten by the upper triangular part of the updated matrix. On entry with *UPLO* = 'L' or 'l', the array  $AP$  must contain the lower triangular part of the Hermitian matrix packed sequentially, column by column, so that  $AP(1)$  contains  $A(1,1)$ ,  $AP(2)$  and  $AP(3)$  contain  $A(2,1)$  and  $A(3,1)$  respectively, and so on. On exit, the array  $AP$  is overwritten by the lower triangular part of the updated matrix. The imaginary parts of the diagonal elements need not be set, they are assumed to be 0, and on exit they are set to 0.

---

## CHER2 or ZHER2 Subroutine

### Purpose

Performs the Hermitian rank 2 operation.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```

SUBROUTINE CHER2(UPLO, N, ALPHA,
X, INCX, Y, INCY, A, LDA)

```

```

COMPLEX ALPHA
INTEGER INCX, INCY, LDA, N
CHARACTER*1 UPLO
COMPLEX A(LDA,*), X(*), Y(*)
SUBROUTINE ZHER2(UPLO, N, ALPHA,
X, INCX, Y, INCY, A, LDA)
COMPLEX*16 ALPHA
INTEGER INCX, INCY, LDA, N
CHARACTER*1 UPLO
COMPLEX*16 A(LDA,*), X(*), Y(*)

```

## Description

The **CHER2** or **ZHER2** subroutine performs the Hermitian rank 2 operation:

$$A := \alpha * x * \text{conjg}(y') + \text{conjg}(\alpha) * y * \text{conjugate}(x') + A$$

where  $\alpha$  is a scalar,  $x$  and  $y$  are  $N$  element vectors and  $A$  is an  $N$  by  $N$  Hermitian matrix.

## Parameters

**UPLO** On entry, *UPLO* specifies whether the upper or lower triangular part of the array *A* is to be referenced as follows:

*UPLO* = 'U' or 'u'  
Only the upper triangular part of *A* is to be referenced.

*UPLO* = 'L' or 'l'  
Only the lower triangular part of *A* is to be referenced.

Unchanged on exit.

**N** On entry, *N* specifies the order of the matrix *A*; *N* must be at least 0; unchanged on exit.

**ALPHA** On entry, *ALPHA* specifies the scalar  $\alpha$ ; unchanged on exit.

**X** A vector of dimension at least  $(1 + (N-1) * \text{abs}(INCX))$ ; on entry, the incremented vector *X* must contain the  $N$  element vector  $x$ ; unchanged on exit.

**INCX** On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must not be 0; unchanged on exit.

**Y** A vector of dimension at least  $(1 + (N-1) * \text{abs}(INCY))$ ; on entry, the incremented vector *Y* must contain the  $N$  element vector  $y$ ; unchanged on exit.

**INCY** On entry, *INCY* specifies the increment for the elements of *Y*; *INCY* must not be 0; unchanged on exit.

**A** An array of dimension  $(LDA, N)$ ; on entry with *UPLO* = 'U' or 'u', the leading  $N$  by  $N$  upper triangular part of the array *A* must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of *A* is not referenced. On exit, the upper triangular part of the array *A* is overwritten by the upper triangular part of the updated matrix. On entry with *UPLO* = 'L' or 'l', the leading  $N$  by  $N$  lower triangular part of the array *A* must contain the lower triangular part of the Hermitian matrix and the strictly upper triangular part of *A* is not referenced. On exit, the lower triangular part of the array *A* is overwritten by the lower triangular part of the updated matrix. The imaginary parts of the diagonal elements need not be set; they are assumed to be 0, and on exit they are set to 0.

**LDA** On entry, *LDA* specifies the first dimension of *A* as declared in the calling (sub) program; *LDA* must be at least  $\max(1, N)$ ; unchanged on exit.

---

## CHPR2 or ZHPR2 Subroutine

### Purpose

Performs the Hermitian rank 2 operation.

### Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

**SUBROUTINE** *CHPR2* (*UPLO*, *N*, *ALPHA*,  
*X*, *INCX*, *Y*, *INCY*, *AP*)

**COMPLEX** *ALPHA*

**INTEGER** *INCX*, *INCY*, *N*

**CHARACTER\*1** *UPLO*

**COMPLEX** *AP*(\*), *X*(\*), *Y*(\*)

**SUBROUTINE**

**ZHPR2**

**COMPLEX\*16** *ALPHA*

**INTEGER** *INCX*, *INCY*, *N*

**CHARACTER\*1** *UPLO*

**COMPLEX\*16** *AP*(\*), *X*(\*), *Y*(\*)

## Description

The **CHPR2** or **ZHPR2** subroutine performs the Hermitian rank 2 operation:

$$A := \alpha * x * \text{conjg}(y') + \text{conjg}(\alpha) * y * \text{conjg}(x') + A$$

where  $\alpha$  is a scalar,  $x$  and  $y$  are  $N$  element vectors and  $A$  is an  $N$  by  $N$  Hermitian matrix, supplied in packed form.

## Parameters

- UPLO* On entry, *UPLO* specifies whether the upper or lower triangular part of the matrix  $A$  is supplied in the packed array *AP* as follows:
- UPLO* = 'U' or 'u'  
The upper triangular part of  $A$  is supplied in *AP*.
- UPLO* = 'L' or 'l'  
The lower triangular part of  $A$  is supplied in *AP*.
- Unchanged on exit.
- N* On entry,  $N$  specifies the order of the matrix  $A$ ;  $N$  must be at least 0; unchanged on exit.
- ALPHA* On entry, *ALPHA* specifies the scalar  $\alpha$ ; unchanged on exit.
- X* A vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; on entry, the incremented array  $X$  must contain the  $N$  element vector  $x$ ; unchanged on exit.
- INCX* On entry, *INCX* specifies the increment for the elements of  $X$ ; *INCX* must not be 0; unchanged on exit.
- Y* A vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCY}))$ ; on entry, the incremented array  $Y$  must contain the  $N$  element vector  $y$ ; unchanged on exit.
- INCY* On entry, *INCY* specifies the increment for the elements of  $Y$ ; *INCY* must not be 0; unchanged on exit.
- AP* A vector of dimension at least  $((N * (N+1)) / 2)$ ; on entry with *UPLO* = 'U' or 'u', the array *AP* must contain the upper triangular part of the Hermitian matrix packed sequentially, column by column, so that *AP*(1) contains  $A(1,1)$ , *AP*(2) and *AP*(3) contain  $A(1,2)$  and  $A(2,2)$  respectively, and so on. On exit, the array *AP* is overwritten by the upper triangular part of the updated matrix. On entry with *UPLO* = 'L' or 'l', the array *AP* must contain the lower triangular part of the Hermitian matrix packed sequentially, column by column, so that *AP*(1) contains  $A(1,1)$ , *AP*(2) and *AP*(3) contain  $A(2,1)$  and  $A(3,1)$  respectively, and so on. On exit, the array *AP* is overwritten by the lower triangular part of the updated matrix. The imaginary parts of the diagonal elements need not be set, they are assumed to be 0, and on exit they are set to 0.

---

## SSYR or DSYR Subroutine

### Purpose

Performs the symmetric rank 1 operation.

## Library

BLAS Library (**libblas.a**)

## FORTTRAN Syntax

```
SUBROUTINE SSYR(UPLO, N, ALPHA,
X, INCX, A, LDA)
REAL ALPHA
INTEGER INCX, LDA, N
CHARACTER*1 UPLO
REAL A(LDA,*), X(*)
SUBROUTINE DSYR(UPLO, N, ALPHA,
X, INCX, A, LDA)
DOUBLE PRECISION ALPHA
INTEGER INCX, LDA, N
CHARACTER*1 UPLO
DOUBLE PRECISION A(LDA,*), X(*)
```

## Description

The **SSYR** or **DSYR** subroutine performs the symmetric rank 1 operation:

$$A := \alpha x x' + A$$

where  $\alpha$  is a real scalar,  $x$  is an  $N$  element vector and  $A$  is an  $N$  by  $N$  symmetric matrix.

## Parameters

**UPLO** On entry, *UPLO* specifies whether the upper or lower triangular part of the array *A* is to be referenced as follows:

**UPLO = 'U' or 'u'**  
Only the upper triangular part of *A* is to be referenced.

**UPLO = 'L' or 'l'**  
Only the lower triangular part of *A* is to be referenced.

Unchanged on exit.

**N** On entry, *N* specifies the order of the matrix *A*; *N* must be at least 0; unchanged on exit.

**ALPHA** On entry, *ALPHA* specifies the scalar  $\alpha$ ; unchanged on exit.

**X** A vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; on entry, the incremented array *X* must contain the  $N$  element vector  $x$ ; unchanged on exit.

**INCX** On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must not be 0; unchanged on exit.

**A** An array of dimension  $(LDA, N)$ ; on entry with *UPLO* = 'U' or 'u', the leading  $N$  by  $N$  upper triangular part of the array *A* must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of *A* is not referenced. On exit, the upper triangular part of the array *A* is overwritten by the upper triangular part of the updated matrix. On entry with *UPLO* = 'L' or 'l', the leading  $N$  by  $N$  lower triangular part of the array *A* must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of *A* is not referenced. On exit, the lower triangular part of the array *A* is overwritten by the lower triangular part of the updated matrix.

**LDA** On entry, *LDA* specifies the first dimension of *A* as declared in the calling (sub) program; *LDA* must be at least  $\max(1, N)$ ; unchanged on exit.

---

## SSPR or DSPR Subroutine

### Purpose

Performs the symmetric rank 1 operation.

## Library

BLAS Library (**libblas.a**)

## FORTTRAN Syntax

```
SUBROUTINE SSPR(UPLO, N, ALPHA,  
X, INCX, AP)
```

```
REAL ALPHA
```

```
INTEGER INCX, N
```

```
CHARACTER*1 UPLO
```

```
REAL AP(*), X(*)
```

```
SUBROUTINE DSPR(UPLO, N, ALPHA,  
X, INCX, AP)
```

```
DOUBLE PRECISION ALPHA
```

```
INTEGER INCX, N
```

```
CHARACTER*1 UPLO
```

```
DOUBLE PRECISION AP(*), X(*)
```

## Description

The **SSPR** or **DSPR** subroutine performs the symmetric rank 1 operation:

$$A := \alpha * x * x' + A$$

where  $\alpha$  is a real scalar,  $x$  is an  $N$  element vector and  $A$  is an  $N$  by  $N$  symmetric matrix, supplied in packed form.

## Parameters

*UPLO* On entry, *UPLO* specifies whether the upper or lower triangular part of the matrix  $A$  is supplied in the packed array *AP* as follows:

*UPLO* = 'U' or 'u'

The upper triangular part of  $A$  is supplied in *AP*.

*UPLO* = 'L' or 'l'

The lower triangular part of  $A$  is supplied in *AP*.

Unchanged on exit.

*N* On entry, *N* specifies the order of the matrix  $A$ ; *N* must be at least 0; unchanged on exit.

*ALPHA* On entry, *ALPHA* specifies the scalar  $\alpha$ ; unchanged on exit.

*X* A vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; on entry, the incremented array *X* must contain the  $N$  element vector  $x$ ; unchanged on exit.

*INCX* On entry, *INCX* specifies the increment for the elements of *X*; *INCX* must not be 0; unchanged on exit.

*AP* A vector of dimension at least  $((N * (N+1)) / 2)$ ; on entry with *UPLO* = 'U' or 'u', the array *AP* must contain the upper triangular part of the symmetric matrix packed sequentially, column by column, so that *AP*(1) contains  $A(1,1)$ , *AP*(2) and *AP*(3) contain  $A(1,2)$  and  $A(2,2)$  respectively, and so on. On exit, the array *AP* is overwritten by the upper triangular part of the updated matrix. On entry with *UPLO* = 'L' or 'l', the array *AP* must contain the lower triangular part of the symmetric matrix packed sequentially, column by column, so that *AP*(1) contains  $A(1,1)$ , *AP*(2) and *AP*(3) contain  $A(2,1)$  and  $A(3,1)$  respectively, and so on. On exit, the array *AP* is overwritten by the lower triangular part of the updated matrix.

---

## SSYR2 or DSYR2 Subroutine

### Purpose

Performs the symmetric rank 2 operation.

## Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

```
SUBROUTINE SSYR2(UPLO, N, ALPHA, X,  
INCX, Y, INCY, A, LDA)  
REAL ALPHA  
INTEGER INCX, INCY, LDA, N  
CHARACTER*1 UPLO  
REAL A(LDA,*), X(*), Y(*)  
SUBROUTINE DSYR2(UPLO, N, ALPHA, X,  
INCX, Y, INCY, A, LDA)  
DOUBLE PRECISION ALPHA  
INTEGER INCX, INCY, LDA, N  
CHARACTER*1 UPLO  
DOUBLE PRECISION A(LDA,*), X(*), Y(*)
```

## Description

The **SSYR2** or **DSYR2** subroutine performs the symmetric rank 2 operation:

$$A := \alpha * x * x' + \alpha * y * y' + A$$

where  $\alpha$  is a scalar,  $x$  and  $y$  are  $N$  element vectors and  $A$  is an  $N$  by  $N$  symmetric matrix.

## Parameters

<i>UPLO</i>	On entry, <i>UPLO</i> specifies whether the upper or lower triangular part of the array <i>A</i> is to be referenced as follows:  <i>UPLO</i> = 'U' or 'u' Only the upper triangular part of <i>A</i> is to be referenced.  <i>UPLO</i> = 'L' or 'l' Only the lower triangular part of <i>A</i> is to be referenced.  Unchanged on exit.
<i>N</i>	On entry, <i>N</i> specifies the order of the matrix <i>A</i> ; <i>N</i> must be at least 0; unchanged on exit.
<i>ALPHA</i>	On entry, <i>ALPHA</i> specifies the scalar $\alpha$ ; unchanged on exit.
<i>X</i>	A vector of dimension at least $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; on entry, the incremented array <i>X</i> must contain the $N$ element vector $x$ ; unchanged on exit.
<i>INCX</i>	On entry, <i>INCX</i> specifies the increment for the elements of <i>X</i> ; <i>INCX</i> must not be 0; unchanged on exit.
<i>Y</i>	A vector of dimension at least $(1 + (N-1) * \text{abs}(\text{INCY}))$ ; on entry, the incremented array <i>Y</i> must contain the $N$ element vector $y$ ; unchanged on exit.
<i>INCY</i>	On entry, <i>INCY</i> specifies the increment for the elements of <i>Y</i> ; <i>INCY</i> must not be 0; unchanged on exit.
<i>A</i>	An array of dimension $(LDA, N)$ ; on entry with <i>UPLO</i> = 'U' or 'u', the leading $N$ by $N$ upper triangular part of the array <i>A</i> must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of <i>A</i> is not referenced. On exit, the upper triangular part of the array <i>A</i> is overwritten by the upper triangular part of the updated matrix. On entry with <i>UPLO</i> = 'L' or 'l', the leading $N$ by $N$ lower triangular part of the array <i>A</i> must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of <i>A</i> is not referenced. On exit, the lower triangular part of the array <i>A</i> is overwritten by the lower triangular part of the updated matrix.
<i>LDA</i>	On entry, <i>LDA</i> specifies the first dimension of <i>A</i> as declared in the calling (sub) program; <i>LDA</i> must be at least $\max(1, N)$ ; unchanged on exit.

---

## SSPR2 or DSPR2 Subroutine

### Purpose

Performs the symmetric rank 2 operation.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE SSPR2(UPLO, N, ALPHA, X,  
INCX, Y, INCY, AP)
```

```
REAL ALPHA
```

```
INTEGER INCX, INCY, N
```

```
CHARACTER*1 UPLO
```

```
REAL AP(*), X(*), Y(*)
```

```
SUBROUTINE DSPR2(UPLO, N, ALPHA, X,  
INCX, Y, INCY, AP)
```

```
DOUBLE PRECISION ALPHA
```

```
INTEGER INCX, INCY, N
```

```
CHARACTER*1 UPLO
```

```
DOUBLE PRECISION AP(*), X(*), Y(*)
```

### Description

The **SSPR2** or **DSPR2** subroutine performs the symmetric rank 2 operation:

$$A := \alpha * x * y' + \alpha * y * x' + A$$

where  $\alpha$  is a scalar,  $x$  and  $y$  are  $N$  element vectors and  $A$  is an  $N$  by  $N$  symmetric matrix, supplied in packed form.

### Parameters

**UPLO** On entry, **UPLO** specifies whether the upper or lower triangular part of the matrix  $A$  is supplied in the packed array  $AP$  as follows:

**UPLO = 'U' or 'u'**

The upper triangular part of  $A$  is supplied in  $AP$ .

**UPLO = 'L' or 'l'**

The lower triangular part of  $A$  is supplied in  $AP$ .

Unchanged on exit.

**N** On entry, **N** specifies the order of the matrix  $A$ ; **N** must be at least 0; unchanged on exit.

**ALPHA** On entry, **ALPHA** specifies the scalar  $\alpha$ ; unchanged on exit.

**X** A vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCX}))$ ; on entry, the incremented array  $X$  must contain the  $N$  element vector  $x$ ; unchanged on exit.

**INCX** On entry, **INCX** specifies the increment for the elements of  $X$ ; **INCX** must not be 0; unchanged on exit.

**Y** A vector of dimension at least  $(1 + (N-1) * \text{abs}(\text{INCY}))$ ; on entry, the incremented array  $Y$  must contain the  $N$  element vector  $y$ ; unchanged on exit.

**INCY** On entry, **INCY** specifies the increment for the elements of  $Y$ ; **INCY** must not be 0; unchanged on exit.

*AP* A vector of dimension at least  $( ( N * (N+1) )/2 )$ ; on entry with *UPLO* = 'U' or 'u', the array *AP* must contain the upper triangular part of the symmetric matrix packed sequentially, column by column, so that *AP*(1) contains *A*(1,1), *AP*(2) and *AP*(3) contain *A*(1,2) and *A*(2,2) respectively, and so on. On exit, the array *AP* is overwritten by the upper triangular part of the updated matrix. On entry with *UPLO* = 'L' or 'l', the array *AP* must contain the lower triangular part of the symmetric matrix packed sequentially, column by column, so that *AP*(1) contains *A*(1,1), *AP*(2) and *AP*(3) contain *A*(2,1) and *A*(3,1) respectively, and so on. On exit, the array *AP* is overwritten by the lower triangular part of the updated matrix.

---

## SGEMM, DGEMM, CGEMM, or ZGEMM Subroutine

### Purpose

Performs matrix-matrix operations on general matrices.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE SGEMM(TRANSA, TRANSB, M, N, K,
ALPHA, A, LDA, B, LDB, BETA, C, LDC)
```

```
CHARACTER*1 TRANSA, TRANSB
```

```
INTEGER M, N, K, LDA, LDB, LDC
```

```
REAL ALPHA, BETA
```

```
REAL A(LDA,*), B(LDB,*), C(LDC,*)
```

```
SUBROUTINE DGEMM(TRANSA, TRANSB, M, N, K,
ALPHA, A, LDA, B, LDB, BETA, C, LDC)
```

```
CHARACTER*1 TRANSA, TRANSB
```

```
INTEGER M, N, K, LDA, LDB, LDC
```

```
DOUBLE PRECISION ALPHA, BETA
```

```
DOUBLE PRECISION A(LDA,*), B(LDB,*), C(LDC,*)
```

```
SUBROUTINE CGEMM(TRANSA, TRANSB, M, N, K,
ALPHA, A, LDA, B, LDB, BETA, C, LDC)
```

```
CHARACTER*1 TRANSA, TRANSB
```

```
INTEGER M, N, K, LDA, LDB, LDC
```

```
COMPLEX ALPHA, BETA
```

```
COMPLEX A(LDA,*), B(LDB,*), C(LDC,*)
```

```
SUBROUTINE ZGEMM(TRANSA, TRANSB, M, N, K,
ALPHA, A, LDA, B, LDB, BETA, C, LDC)
```

```
CHARACTER*1 TRANSA, TRANSB
```

```
INTEGER M, N, K, LDA, LDB, LDC
```

```
COMPLEX*16 ALPHA, BETA
```

```
COMPLEX*16 A(LDA,*), B(LDB,*), C(LDC,*)
```

### Description

The **SGEMM**, **DGEMM**, **CGEMM**, or **ZGEMM** subroutine performs one of the matrix-matrix operations:

$$C := \alpha * \text{op}(A) * \text{op}(B) + \beta * C$$

where *op*(*X*) is one of *op*(*X*) = *X* or *op*(*X*) = *X*' ,  $\alpha$  and  $\beta$  are scalars, and *A*, *B* and *C* are matrices, with *op*(*A*) an *M* by *K* matrix, *op*(*B*) a *K* by *N* matrix and *C* an *M* by *N* matrix.

## Parameters

- TRANSA* On entry, *TRANSA* specifies the form of  $op(A)$  to be used in the matrix multiplication as follows:
- TRANSA* = 'N' or 'n'  
 $op(A) = A$
- TRANSA* = 'T' or 't'  
 $op(A) = A'$
- TRANSA* = 'C' or 'c'  
 $op(A) = A'$
- Unchanged on exit.
- TRANSB* On entry, *TRANSB* specifies the form of  $op(B)$  to be used in the matrix multiplication as follows:
- TRANSB* = 'N' or 'n'  
 $op(B) = B$
- TRANSB* = 'T' or 't'  
 $op(B) = B'$
- TRANSB* = 'C' or 'c'  
 $op(B) = B'$
- Unchanged on exit.
- M* On entry, *M* specifies the number of rows of the matrix  $op(A)$  and of the matrix *C*; *M* must be at least 0; unchanged on exit.
- N* On entry, *N* specifies the number of columns of the matrix  $op(B)$  and the number of columns of the matrix *C*; *N* must be at least 0; unchanged on exit.
- K* On entry, *K* specifies the number of columns of the matrix  $op(A)$  and the number of rows of the matrix  $op(B)$ ; *K* must be at least 0; unchanged on exit.
- ALPHA* On entry, *ALPHA* specifies the scalar alpha; unchanged on exit.
- A* An array of dimension ( *LDA*, *KA* ), where *KA* is *K* when *TRANSA* = 'N' or 'n', and is *M* otherwise; on entry with *TRANSA* = 'N' or 'n', the leading *M* by *K* part of the array *A* must contain the matrix *A*, otherwise the leading *K* by *M* part of the array *A* must contain the matrix *A*; unchanged on exit.
- LDA* On entry, *LDA* specifies the first dimension of *A* as declared in the calling (sub) program. When *TRANSA* = 'N' or 'n' then *LDA* must be at least  $\max(1, M)$ , otherwise *LDA* must be at least  $\max(1, K)$ ; unchanged on exit.
- B* An array of dimension ( *LDB*, *KB* ) where *KB* is *N* when *TRANSB* = 'N' or 'n', and is *K* otherwise; on entry with *TRANSB* = 'N' or 'n', the leading *K* by *N* part of the array *B* must contain the matrix *B*, otherwise the leading *N* by *K* part of the array *B* must contain the matrix *B*; unchanged on exit.
- LDB* On entry, *LDB* specifies the first dimension of *B* as declared in the calling (sub) program. When *TRANSB* = 'N' or 'n' then *LDB* must be at least  $\max(1, K)$ , otherwise *LDB* must be at least  $\max(1, N)$ ; unchanged on exit.
- BETA* On entry, *BETA* specifies the scalar beta. When *BETA* is supplied as 0 then *C* need not be set on input; unchanged on exit.
- C* An array of dimension ( *LDC*, *N* ); on entry, the leading *M* by *N* part of the array *C* must contain the matrix *C*, except when beta is 0, in which case *C* need not be set on entry; on exit, the array *C* is overwritten by the *M* by *N* matrix (  $\alpha * op(A) * op(B) + \beta * C$  ).
- LDC* On entry, *LDC* specifies the first dimension of *C* as declared in the calling (sub) program; *LDC* must be at least  $\max(1, M)$ ; unchanged on exit.

---

## SSYMM, DSYMM, CSYMM, or ZSYMM Subroutine

### Purpose

Performs matrix-matrix matrix operations on symmetric matrices.

## Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

```
SUBROUTINE SSYMM(SIDE, UPLO, M, N, ALPHA,  
A, LDA, B, LDB, BETA, C, LDC)
```

```
CHARACTER*1 SIDE, UPLO
```

```
INTEGER M, N, LDA, LDB, LDC
```

```
REAL ALPHA, BETA
```

```
REAL A(LDA,*), B(LDB,*), C(LDC,*)
```

```
SUBROUTINE DSYMM(SIDE, UPLO, M, N, ALPHA,  
A, LDA, B, LDB, BETA, C, LDC)
```

```
CHARACTER*1 SIDE, UPLO
```

```
INTEGER M, N, LDA, LDB, LDC
```

```
DOUBLE PRECISION ALPHA, BETA
```

```
DOUBLE PRECISION A(LDA,*), B(LDB,*), C(LDC,*)
```

```
SUBROUTINE CSYMM(SIDE, UPLO, M, N, ALPHA,  
A, LDA, B, LDB, BETA, C, LDC)
```

```
CHARACTER*1 SIDE, UPLO
```

```
INTEGER M, N, LDA, LDB, LDC
```

```
COMPLEX ALPHA, BETA
```

```
COMPLEX A(LDA,*), B(LDB,*), C(LDC,*)
```

```
SUBROUTINE ZSYMM(SIDE, UPLO, M, N, ALPHA,  
A, LDA, B, LDB, BETA, C, LDC)
```

```
CHARACTER*1 SIDE, UPLO
```

```
INTEGER M, N, LDA, LDB, LDC
```

```
COMPLEX*16 ALPHA, BETA
```

```
COMPLEX*16 A(LDA,*), B(LDB,*), C(LDC,*)
```

## Description

The **SSYMM**, **DSYMM**, **CSYMM**, or **ZSYMM** subroutine performs one of the matrix-matrix operations:

$$C := \alpha * A * B + \beta * C$$

OR

$$C := \alpha * B * A + \beta * C$$

where  $\alpha$  and  $\beta$  are scalars,  $A$  is a symmetric matrix and  $B$  and  $C$  are  $M$  by  $N$  matrices.

## Parameters

*SIDE* On entry, *SIDE* specifies whether the symmetric matrix  $A$  appears on the left or right in the operation as follows:

*SIDE* = 'L' or 'l'

$$C := \alpha * A * B + \beta * C$$

*SIDE* = 'R' or 'r'

$$C := \alpha * B * A + \beta * C$$

Unchanged on exit.

*UPLO* On entry, *UPLO* specifies whether the upper or lower triangular part of the symmetric matrix *A* is to be referenced as follows:

*UPLO* = 'U' or 'u'  
Only the upper triangular part of the symmetric matrix is to be referenced.

*UPLO* = 'L' or 'l'  
Only the lower triangular part of the symmetric matrix is to be referenced.

Unchanged on exit.

*M* On entry, *M* specifies the number of rows of the matrix *C*; *M* must be at least 0; unchanged on exit.

*N* On entry, *N* specifies the number of columns of the matrix *C*; *N* must be at least 0; unchanged on exit.

*ALPHA* On entry, *ALPHA* specifies the scalar alpha; unchanged on exit.

*A* An array of dimension ( *LDA*, *KA* ), where *KA* is *M* when *SIDE* = 'L' or 'l' and is *N* otherwise; on entry with *SIDE* = 'L' or 'l', the *M* by *M* part of the array *A* must contain the symmetric matrix, such that when *UPLO* = 'U' or 'u', the leading *M* by *M* upper triangular part of the array *A* must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of *A* is not referenced, and when *UPLO* = 'L' or 'l', the leading *M* by *M* lower triangular part of the array *A* must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of *A* is not referenced. On entry with *SIDE* = 'R' or 'r', the *N* by *N* part of the array *A* must contain the symmetric matrix, such that when *UPLO* = 'U' or 'u', the leading *N* by *N* upper triangular part of the array *A* must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of *A* is not referenced, and when *UPLO* = 'L' or 'l', the leading *N* by *N* lower triangular part of the array *A* must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of *A* is not referenced; unchanged on exit.

*LDA* On entry, *LDA* specifies the first dimension of *A* as declared in the calling (sub) program. When *SIDE* = 'L' or 'l' then *LDA* must be at least max( 1, *M* ), otherwise *LDA* must be at least max( 1, *N* ); unchanged on exit.

*B* An array of dimension ( *LDB*, *N* ); on entry, the leading *M* by *N* part of the array *B* must contain the matrix *B*; unchanged on exit.

*LDB* On entry, *LDB* specifies the first dimension of *B* as declared in the calling (sub) program; *LDB* must be at least max( 1, *M* ); unchanged on exit.

*BETA* On entry, *BETA* specifies the scalar beta; when *BETA* is supplied as 0 then *C* need not be set on input; unchanged on exit.

*C* An array of dimension ( *LDC*, *N* ); on entry, the leading *M* by *N* part of the array *C* must contain the matrix *C*, except when beta is 0, in which case *C* need not be set on entry; on exit, the array *C* is overwritten by the *M* by *N* updated matrix.

*LDC* On entry, *LDC* specifies the first dimension of *C* as declared in the calling (sub) program; *LDC* must be at least max( 1, *M* ); unchanged on exit.

---

## CHEMM or ZHEMM Subroutine

### Purpose

Performs matrix-matrix operations on Hermitian matrices.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```

SUBROUTINE CHEMM(SIDE, UPLO, M, N, ALPHA, A,
LDA, B, LDB, BETA, C, LDC)
CHARACTER*1 SIDE, UPLO
INTEGER M, N, LDA, LDB, LDC
COMPLEX ALPHA, BETA
COMPLEX A(LDA,*), B(LDB,*), C(LDC,*)

```

```

SUBROUTINE ZHEMM(SIDE, UPLO, M, N, ALPHA, A,
LDA, B, LDB, BETA, C, LDC)
CHARACTER*1 SIDE,UPLO
INTEGER M,N,LDA,LDB,LDC
COMPLEX*16 ALPHA,BETA
COMPLEX*16 A(LDA,*), B(LDB,*), C(LDC,*)

```

## Purpose

The **CHEMM** or **ZHEMM** subroutine performs one of the matrix-matrix operations:

$C := \alpha * A * B + \beta * C$

OR

$C := \alpha * B * A + \beta * C$

where  $\alpha$  and  $\beta$  are scalars,  $A$  is an Hermitian matrix, and  $B$  and  $C$  are  $M$  by  $N$  matrices.

## Parameters

**SIDE** On entry, *SIDE* specifies whether the Hermitian matrix  $A$  appears on the left or right in the operation as follows:

*SIDE* = 'L' or 'l'

$C := \alpha * A * B + \beta * C$

*SIDE* = 'R' or 'r'

$C := \alpha * B * A + \beta * C$

Unchanged on exit.

**UPLO** On entry, *UPLO* specifies whether the upper or lower triangular part of the Hermitian matrix  $A$  is to be referenced as follows:

*UPLO* = 'U' or 'u'

Only the upper triangular part of the Hermitian matrix is to be referenced.

*UPLO* = 'L' or 'l'

Only the lower triangular part of the Hermitian matrix is to be referenced.

Unchanged on exit.

**M** On entry,  $M$  specifies the number of rows of the matrix  $C$ ;  $M$  must be at least 0; unchanged on exit.

**N** On entry,  $N$  specifies the number of columns of the matrix  $C$ ;  $N$  must be at least 0; unchanged on exit.

**ALPHA** On entry, *ALPHA* specifies the scalar  $\alpha$ ; unchanged on exit.

**A** An array of dimension ( *LDA*, *KA* ), where *KA* is  $M$  when *SIDE* = 'L' or 'l' and is  $N$  otherwise; on entry with *SIDE* = 'L' or 'l', the  $M$  by  $M$  part of the array  $A$  must contain the Hermitian matrix, such that when *UPLO* = 'U' or 'u', the leading  $M$  by  $M$  upper triangular part of the array  $A$  must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of  $A$  is not referenced, and when *UPLO* = 'L' or 'l', the leading  $M$  by  $M$  lower triangular part of the array  $A$  must contain the lower triangular part of the Hermitian matrix and the strictly upper triangular part of  $A$  is not referenced; on entry with *SIDE* = 'R' or 'r', the  $N$  by  $N$  part of the array  $A$  must contain the Hermitian matrix, such that when *UPLO* = 'U' or 'u', the leading  $N$  by  $N$  upper triangular part of the array  $A$  must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of  $A$  is not referenced, and when *UPLO* = 'L' or 'l', the leading  $N$  by  $N$  lower triangular part of the array  $A$  must contain the lower triangular part of the Hermitian matrix and the strictly upper triangular part of  $A$  is not referenced. The imaginary parts of the diagonal elements need not be set, they are assumed to be 0; unchanged on exit.

**LDA** On entry, *LDA* specifies the first dimension of  $A$  as declared in the calling (sub) program. When *SIDE* = 'L' or 'l' then *LDA* must be at least  $\max(1, M)$ , otherwise *LDA* must be at least  $\max(1, N)$ ; unchanged on exit.

**B** An array of dimension ( *LDB*,  $N$  ); on entry, the leading  $M$  by  $N$  part of the array  $B$  must contain the matrix  $B$ ; unchanged on exit.

**LDB** On entry, *LDB* specifies the first dimension of  $B$  as declared in the calling (sub) program; *LDB* must be at least  $\max(1, M)$ ; unchanged on exit.

<i>BETA</i>	On entry, <i>BETA</i> specifies the scalar beta. When <i>BETA</i> is supplied as 0 then <i>C</i> need not be set on input; unchanged on exit.
<i>C</i>	An array of dimension ( <i>LDC</i> , <i>N</i> ); on entry, the leading <i>M</i> by <i>N</i> part of the array <i>C</i> must contain the matrix <i>C</i> , except when beta is 0, in which case <i>C</i> need not be set on entry; on exit, the array <i>C</i> is overwritten by the <i>M</i> by <i>N</i> updated matrix.
<i>LDC</i>	On entry, <i>LDC</i> specifies the first dimension of <i>C</i> as declared in the calling (sub) program; <i>LDC</i> must be at least $\max(1, M)$ ; unchanged on exit.

---

## SSYRK, DSYRK, CSYRK, or ZSYRK Subroutine

### Purpose

Perform symmetric rank k operations.

### Library

BLAS Library (**libblas.a**)

### FORTTRAN Syntax

```
SUBROUTINE SSYRK(UPLO, TRANS, N, K, ALPHA,
A, LDA, BETA, C, LDC)
```

```
CHARACTER*1 UPLO, TRANS
```

```
INTEGER N, K, LDA, LDC
```

```
REAL ALPHA, BETA
```

```
REAL A(LDA,*), C(LDC,*)
```

```
SUBROUTINE DSYRK(UPLO, TRANS, N, K, ALPHA,
A, LDA, BETA, C, LDC)
```

```
CHARACTER*1 UPLO, TRANS
```

```
INTEGER N, K, LDA, LDC
```

```
DOUBLE PRECISION ALPHA, BETA
```

```
DOUBLE PRECISION A(LDA,*), C(LDC,*)
```

```
SUBROUTINE CSYRK(UPLO, TRANS, N, K, ALPHA,
A, LDA, BETA, C, LDC)
```

```
CHARACTER*1 UPLO, TRANS
```

```
INTEGER N, K, LDA, LDC
```

```
COMPLEX ALPHA, BETA
```

```
COMPLEX A(LDA,*), C(LDC,*)
```

```
SUBROUTINE ZSYRK(UPLO, TRANS, N, K, ALPHA,
A, LDA, BETA, C, LDC)
```

```
CHARACTER*1 UPLO, TRANS
```

```
INTEGER N, K, LDA, LDC
```

```
COMPLEX*16 ALPHA, BETA
```

```
COMPLEX*16 A(LDA,*), C(LDC,*)
```

### Description

The **SSYRK**, **DSYRK**, **CSYRK** or **ZSYRK** subroutine performs one of the symmetric rank k operations:

$$C := \alpha * A * A' + \beta * C$$

OR

$$C := \alpha * A' * A + \beta * C$$

where alpha and beta are scalars, *C* is an *N* by *N* symmetric matrix, and *A* is an *N* by *K* matrix in the first case and a *K* by *N* matrix in the second case.

## Parameters

- UPLO* On entry, *UPLO* specifies whether the upper or lower triangular part of the array *C* is to be referenced as follows:
- UPLO* = 'U' or 'u'  
Only the upper triangular part of *C* is to be referenced.
- UPLO* = 'L' or 'l'  
Only the lower triangular part of *C* is to be referenced.
- Unchanged on exit.
- TRANS* On entry, *TRANS* specifies the operation to be performed as follows:
- TRANS* = 'N' or 'n'  
 $C := \alpha * A * A' + \beta * C$
- TRANS* = 'T' or 't'  
 $C := \alpha * A' * A + \beta * C$
- TRANS* = 'C' or 'c'  
 $C := \alpha * A' * A + \beta * C$
- Unchanged on exit.
- N* On entry, *N* specifies the order of the matrix *C*; *N* must be at least 0; unchanged on exit.
- K* On entry with *TRANS* = 'N' or 'n', *K* specifies the number of columns of the matrix *A*, and on entry with *TRANS* = 'T' or 't' or 'C' or 'c', *K* specifies the number of rows of the matrix *A*; *K* must be at least 0; unchanged on exit.
- ALPHA* On entry, *ALPHA* specifies the scalar alpha; unchanged on exit.
- A* An array of dimension ( *LDA*, *KA* ), where *KA* is *K* when *TRANS* = 'N' or 'n', and is *N* otherwise; on entry with *TRANS* = 'N' or 'n', the leading *N* by *K* part of the array *A* must contain the matrix *A*, otherwise the leading *K* by *N* part of the array *A* must contain the matrix *A*; unchanged on exit.
- LDA* On entry, *LDA* specifies the first dimension of *A* as declared in the calling (sub) program. When *TRANS* = 'N' or 'n', *LDA* must be at least max( 1, *N* ); otherwise *LDA* must be at least max( 1, *K* ); unchanged on exit.
- BETA* On entry, *BETA* specifies the scalar beta; unchanged on exit.
- C* An array of dimension ( *LDC*, *N* ); on entry with *UPLO* = 'U' or 'u', the leading *N* by *N* upper triangular part of the array *C* must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of *C* is not referenced; on exit, the upper triangular part of the array *C* is overwritten by the upper triangular part of the updated matrix; on entry with *UPLO* = 'L' or 'l', the leading *N* by *N* lower triangular part of the array *C* must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of *C* is not referenced; on exit, the lower triangular part of the array *C* is overwritten by the lower triangular part of the updated matrix.
- LDC* On entry, *LDC* specifies the first dimension of *C* as declared in the calling (sub) program; *LDC* must be at least max( 1, *N* ); unchanged on exit.

---

## CHERK or ZHERK Subroutine

### Purpose

Performs Hermitian rank k operations.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

**SUBROUTINE** **CHERK**(*UPLO*, *TRANS*, *N*, *K*, *ALPHA*,  
*A*, *LDA*, *BETA*, *C*, *LDC*)  
**CHARACTER\***1 *UPLO*, *TRANS*

```

INTEGER N, K, LDA, LDC
REAL ALPHA, BETA
COMPLEX A(LDA,*), C(LDC,*)
SUBROUTINE ZHERK(UPLO, TRANS, N, K, ALPHA,
A, LDA, BETA, C, LDC)
CHARACTER*1 UPLO, TRANS
INTEGER N, K, LDA, LDC
DOUBLE PRECISION ALPHA, BETA
COMPLEX*16 A(LDA,*), C(LDC,*)

```

## Description

The **CHERK** or **ZHERK** subroutine performs one of the Hermitian rank k operations:

$$C := \alpha * A * \text{conjg}(A') + \beta * C$$

OR

$$C := \alpha * \text{conjg}(A') * A + \beta * C$$

where  $\alpha$  and  $\beta$  are real scalars,  $C$  is an  $N$  by  $N$  Hermitian matrix, and  $A$  is an  $N$  by  $K$  matrix in the first case and a  $K$  by  $N$  matrix in the second case.

## Parameters

**UPLO** On entry, **UPLO** specifies whether the upper or lower triangular part of the array  $C$  is to be referenced as follows:

**UPLO = 'U' or 'u'**

Only the upper triangular part of  $C$  is to be referenced.

**UPLO = 'L' or 'l'**

Only the lower triangular part of  $C$  is to be referenced.

Unchanged on exit.

**TRANS** On entry, **TRANS** specifies the operation to be performed as follows:

**TRANS = 'N' or 'n'**

$$C := \alpha * A * \text{conjg}(A') + \beta * C$$

**TRANS = 'C' or 'c'**

$$C := \alpha * \text{conjg}(A') * A + \beta * C$$

Unchanged on exit.

**N** On entry, **N** specifies the order of the matrix  $C$ ; **N** must be at least 0; unchanged on exit.

**K** On entry with **TRANS = 'N' or 'n'**, **K** specifies the number of columns of the matrix  $A$ , and on entry with **TRANS = 'C' or 'c'**, **K** specifies the number of rows of the matrix  $A$ ; **K** must be at least 0; unchanged on exit.

**ALPHA** On entry, **ALPHA** specifies the scalar  $\alpha$ ; unchanged on exit.

**A** An array of dimension ( **LDA**, **KA** ), where **KA** is **K** when **TRANS = 'N' or 'n'**, and is **N** otherwise; on entry with **TRANS = 'N' or 'n'**, the leading **N** by **K** part of the array  $A$  must contain the matrix  $A$ , otherwise the leading **K** by **N** part of the array  $A$  must contain the matrix  $A$ ; unchanged on exit.

**LDA** On entry, **LDA** specifies the first dimension of  $A$  as declared in the calling (sub) program. When **TRANS = 'N' or 'n'**, **LDA** must be at least  $\max(1, N)$ , otherwise **LDA** must be at least  $\max(1, K)$ ; unchanged on exit.

**BETA** On entry, **BETA** specifies the scalar  $\beta$ ; unchanged on exit.

- C* An array of dimension ( *LDC*, *N* ); on entry with *UPLO* = 'U' or 'u', the leading *N* by *N* upper triangular part of the array *C* must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of *C* is not referenced; on exit, the upper triangular part of the array *C* is overwritten by the upper triangular part of the updated matrix; on entry with *UPLO* = 'L' or 'l', the leading *N* by *N* lower triangular part of the array *C* must contain the lower triangular part of the Hermitian matrix and the strictly upper triangular part of *C* is not referenced; on exit, the lower triangular part of the array *C* is overwritten by the lower triangular part of the updated matrix. The imaginary parts of the diagonal elements need not be set, they are assumed to be 0, and on exit they are set to 0.
- LDC* On entry, *LDC* specifies the first dimension of *C* as declared in the calling (sub) program; *LDC* must be at least max( 1, *N* ); unchanged on exit.

## SSYR2K, DSYR2K, CSYR2K, or ZSYR2K Subroutine

### Purpose

Performs symmetric rank 2k operations.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE SSYR2K(UPLO, TRANS, N, K, ALPHA,  
A, LDA, B, LDB, BETA, C, LDC)
```

```
CHARACTER*1 UPLO, TRANS
```

```
INTEGER N, K, LDA, LDB, LDC
```

```
REAL ALPHA, BETA
```

```
REAL A(LDA,*), B(LDB,*), C(LDC,*)
```

```
SUBROUTINE DSYR2K(UPLO, TRANS, N, K, ALPHA,  
A, LDA, B, LDB, BETA, C, LDC)
```

```
CHARACTER*1 UPLO, TRANS
```

```
INTEGER N, K, LDA, LDB, LDC
```

```
DOUBLE PRECISION ALPHA, BETA
```

```
DOUBLE PRECISION A(LDA,*), B(LDB,*), C(LDC,*)
```

```
SUBROUTINE CSYR2K(UPLO, TRANS, N, K, ALPHA,  
A, LDA, B, LDB, BETA, C, LDC)
```

```
CHARACTER*1 UPLO, TRANS
```

```
INTEGER N, K, LDA, LDB, LDC
```

```
COMPLEX ALPHA, BETA
```

```
COMPLEX A(LDA,*), B(LDB,*), C(LDC,*)
```

```
SUBROUTINE ZSYR2K(UPLO, TRANS, N, K, ALPHA,  
A, LDA, B, LDB, BETA, C, LDC)
```

```
CHARACTER*1 UPLO, TRANS
```

```
INTEGER N, K, LDA, LDB, LDC
```

```
COMPLEX*16 ALPHA, BETA
```

```
COMPLEX*16 A(LDA,*), B(LDB,*), C(LDC,*)
```

### Description

The **SSYR2K**, **DSYR2K**, **CSYR2K**, or **ZSYR2K** subroutine performs one of the symmetric rank 2k operations:

$$C := \alpha * A * B' + \alpha * B * A' + \beta * C$$

OR

$$C := \alpha * A' * B + \alpha * B' * A + \beta * C$$

where alpha and beta are scalars,  $C$  is an  $N$  by  $N$  symmetric matrix, and  $A$  and  $B$  are  $N$  by  $K$  matrices in the first case and  $K$  by  $N$  matrices in the second case.

## Parameters

- UPLO** On entry, *UPLO* specifies whether the upper or lower triangular part of the array  $C$  is to be referenced as follows:
- UPLO = 'U' or 'u'**  
Only the upper triangular part of  $C$  is to be referenced.
- UPLO = 'L' or 'l'**  
Only the lower triangular part of  $C$  is to be referenced.
- Unchanged on exit.
- TRANS** On entry, *TRANS* specifies the operation to be performed as follows:
- TRANS = 'N' or 'n'**  
 $C := \alpha * A * B + \alpha * B * A + \beta * C$
- TRANS = 'T' or 't'**  
 $C := \alpha * A' * B + \alpha * B' * A + \beta * C$
- Unchanged on exit.
- N** On entry,  $N$  specifies the order of the matrix  $C$ ;  $N$  must be at least 0; unchanged on exit.
- K** On entry with *TRANS* = 'N' or 'n',  $K$  specifies the number of columns of the matrices  $A$  and  $B$ , and on entry with *TRANS* = 'T' or 't',  $K$  specifies the number of rows of the matrices  $A$  and  $B$ ;  $K$  must be at least 0; unchanged on exit.
- ALPHA** On entry, *ALPHA* specifies the scalar alpha; unchanged on exit.
- A** An array of dimension ( *LDA*, *KA* ), where *KA* is  $K$  when *TRANS* = 'N' or 'n', and is  $N$  otherwise; on entry with *TRANS* = 'N' or 'n', the leading  $N$  by  $K$  part of the array  $A$  must contain the matrix  $A$ , otherwise the leading  $K$  by  $N$  part of the array  $A$  must contain the matrix  $A$ ; unchanged on exit.
- LDA** On entry, *LDA* specifies the first dimension of  $A$  as declared in the calling (sub) program. When *TRANS* = 'N' or 'n', *LDA* must be at least  $\max(1, N)$ ; otherwise *LDA* must be at least  $\max(1, K)$ ; unchanged on exit.
- B** An array of dimension ( *LDB*, *KB* ), where *KB* is  $K$  when *TRANS* = 'N' or 'n', and is  $N$  otherwise; on entry with *TRANS* = 'N' or 'n', the leading  $N$  by  $K$  part of the array  $B$  must contain the matrix  $B$ , otherwise the leading  $K$  by  $N$  part of the array  $B$  must contain the matrix  $B$ ; unchanged on exit.
- LDB** On entry, *LDB* specifies the first dimension of  $B$  as declared in the calling (sub) program. When *TRANS* = 'N' or 'n', *LDB* must be at least  $\max(1, N)$ ; otherwise *LDB* must be at least  $\max(1, K)$ ; unchanged on exit.
- BETA** On entry, *BETA* specifies the scalar beta; unchanged on exit.
- C** An array of dimension ( *LDC*,  $N$  ); on entry with *UPLO* = 'U' or 'u', the leading  $N$  by  $N$  upper triangular part of the array  $C$  must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of  $C$  is not referenced; on exit, the upper triangular part of the array  $C$  is overwritten by the upper triangular part of the updated matrix. On entry with *UPLO* = 'L' or 'l', the leading  $N$  by  $N$  lower triangular part of the array  $C$  must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of  $C$  is not referenced; on exit, the lower triangular part of the array  $C$  is overwritten by the lower triangular part of the updated matrix.
- LDC** On entry, *LDC* specifies the first dimension of  $C$  as declared in the calling (sub) program; *LDC* must be at least  $\max(1, N)$ ; unchanged on exit.

---

## CHER2K or ZHER2K Subroutine

### Purpose

Performs Hermitian rank 2k operations.

## Library

BLAS Library (**libblas.a**)

## FORTRAN Syntax

**SUBROUTINE** **CHER2K**(*UPLO*, *TRANS*, *N*, *K*, *ALPHA*,  
*A*, *LDA*, *B*, *LDB*, *C*, *LDC*)

**CHARACTER\*1** *UPLO*, *TRANS*

**INTEGER** *N*, *K*, *LDA*, *LDB*, *LDC*

**REAL** *BETA*

**COMPLEX** *ALPHA*

**COMPLEX** *A*(*LDA*,\*), *B*(*LDB*,\*), *C*(*LDC*,\*)

**SUBROUTINE** **ZHER2K**(*UPLO*, *TRANS*, *N*, *K*, *ALPHA*,  
*A*, *LDA*, *B*, *LDB*, *C*, *LDC*)

**CHARACTER\*1** *UPLO*, *TRANS*

**INTEGER** *N*, *K*, *LDA*, *LDB*, *LDC*

**DOUBLE PRECISION** *BETA*

**COMPLEX\*16** *ALPHA*

**COMPLEX\*16** *A*(*LDA*,\*), *B*(*LDB*,\*), *C*(*LDC*,\*)

## Description

The **CHER2K** or **ZHER2K** subroutine performs one of the Hermitian rank 2k operations:

$C := \alpha * A * \text{conjg}(B') + \text{conjg}(\alpha) * B * \text{conjg}(A') + \beta * C$

OR

$C := \alpha * \text{conjg}(A') * B + \text{conjg}(\alpha) * \text{conjg}(B') * A + \beta * C$

where  $\alpha$  and  $\beta$  are scalars with  $\beta$  real,  $C$  is an  $N$  by  $N$  Hermitian matrix, and  $A$  and  $B$  are  $N$  by  $K$  matrices in the first case and  $K$  by  $N$  matrices in the second case.

## Parameters

**UPLO** On entry, *UPLO* specifies whether the upper or lower triangular part of the array  $C$  is to be referenced as follows:

**UPLO = 'U' or 'u'**

Only the upper triangular part of  $C$  is to be referenced.

**UPLO = 'L' or 'l'**

Only the lower triangular part of  $C$  is to be referenced.

Unchanged on exit.

**TRANS** On entry, *TRANS* specifies the operation to be performed as follows:

**TRANS = 'N' or 'n'**

$C := \alpha * A * \text{conjg}(B') + \text{conjg}(\alpha) * B * \text{conjg}(A') + \beta * C$

**TRANS = 'C' or 'c'**

$C := \alpha * \text{conjg}(A') * B + \text{conjg}(\alpha) * \text{conjg}(B') * A + \beta * C$

Unchanged on exit.

**N** On entry,  $N$  specifies the order of the matrix  $C$ ;  $N$  must be at least 0; unchanged on exit.

**K** On entry with *TRANS* = 'N' or 'n',  $K$  specifies the number of columns of the matrices  $A$  and  $B$ , and on entry with *TRANS* = 'C' or 'c',  $K$  specifies the number of rows of the matrices  $A$  and  $B$ ;  $K$  must be at least 0; unchanged on exit.

**ALPHA** On entry, *ALPHA* specifies the scalar  $\alpha$ ; unchanged on exit.

<i>A</i>	An array of dimension ( <i>LDA</i> , <i>KA</i> ), where <i>KA</i> is <i>K</i> when <i>TRANS</i> = 'N' or 'n', and is <i>N</i> otherwise; on entry with <i>TRANS</i> = 'N' or 'n', the leading <i>N</i> by <i>K</i> part of the array <i>A</i> must contain the matrix <i>A</i> , otherwise the leading <i>K</i> by <i>N</i> part of the array <i>A</i> must contain the matrix <i>A</i> ; unchanged on exit.
<i>LDA</i>	On entry, <i>LDA</i> specifies the first dimension of <i>A</i> as declared in the calling (sub) program. When <i>TRANS</i> = 'N' or 'n', <i>LDA</i> must be at least max( 1, <i>N</i> ); otherwise <i>LDA</i> must be at least max( 1, <i>K</i> ); unchanged on exit.
<i>B</i>	An array of dimension ( <i>LDB</i> , <i>KB</i> ), where <i>KB</i> is <i>K</i> when <i>TRANS</i> = 'N' or 'n', and is <i>N</i> otherwise; on entry with <i>TRANS</i> = 'N' or 'n', the leading <i>N</i> by <i>K</i> part of the array <i>B</i> must contain the matrix <i>B</i> , otherwise the leading <i>K</i> by <i>N</i> part of the array <i>B</i> must contain the matrix <i>B</i> ; unchanged on exit.
<i>LDB</i>	On entry, <i>LDB</i> specifies the first dimension of <i>B</i> as declared in the calling (sub) program. When <i>TRANS</i> = 'N' or 'n', <i>LDB</i> must be at least max( 1, <i>N</i> ); otherwise <i>LDB</i> must be at least max( 1, <i>K</i> ); unchanged on exit.
<i>BETA</i>	On entry, <i>BETA</i> specifies the scalar beta; unchanged on exit.
<i>C</i>	An array of dimension ( <i>LDC</i> , <i>N</i> ); on entry with <i>UPLO</i> = 'U' or 'u', the leading <i>N</i> by <i>N</i> upper triangular part of the array <i>C</i> must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of <i>C</i> is not reference; on exit, the upper triangular part of the array <i>C</i> is overwritten by the upper triangular part of the updated matrix; on entry with <i>UPLO</i> = 'L' or 'l', the leading <i>N</i> by <i>N</i> lower triangular part of the array <i>C</i> must contain the lower triangular part of the Hermitian matrix and the strictly upper triangular part of <i>C</i> is not referenced; on exit, the lower triangular part of the array <i>C</i> is overwritten by the lower triangular part of the updated matrix. The imaginary parts of the diagonal elements need not be set, they are assumed to be 0, and on exit they are set to 0.
<i>LDC</i>	On entry, <i>LDC</i> specifies the first dimension of <i>C</i> as declared in the calling (sub) program; <i>LDC</i> must be at least max( 1, <i>N</i> ); unchanged on exit.

---

## STRMM, DTRMM, CTRMM, or ZTRMM Subroutine

### Purpose

Performs matrix-matrix operations on triangular matrices.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE STRMM(SIDE, UPLO, TRANS, DIAG,  
M, N, ALPHA, A, LDA, B, LDB)
```

```
CHARACTER*1 SIDE, UPLO, TRANS, DIAG
```

```
INTEGER M, N, LDA, LDB
```

```
REAL ALPHA
```

```
REAL A(LDA,*), B(LDB,*)
```

```
SUBROUTINE DTRMM(SIDE, UPLO, TRANS, DIAG,  
M, N, ALPHA, A, LDA, B, LDB)
```

```
CHARACTER*1
```

```
SIDE,UPLO,TRANS,DIAG
```

```
INTEGER M,N,LDA,LDB
```

```
DOUBLE PRECISION ALPHA
```

```
DOUBLE PRECISION A(LDA,*), B(LDB,*)
```

```
SUBROUTINE CTRMM(SIDE, UPLO, TRANS, DIAG,  
M, N, ALPHA, A, LDA, B, LDB)
```

```
CHARACTER*1
```

```
SIDE,UPLO,TRANS,DIAG
```

```
INTEGER M,N,LDA,LDB
```

```
COMPLEX ALPHA
```

```
COMPLEX A(LDA,*), B(LDB,*)
```

```
SUBROUTINE ZTRMM(SIDE, UPLO, TRANS, DIAG,  
M, N, ALPHA, A, LDA, B, LDB)
```

```
CHARACTER*1
```

*SIDE, UPLO, TRANSA, DIAG*  
**INTEGER** *M, N, LDA, LDB*  
**COMPLEX\*16** *ALPHA*  
**COMPLEX\*16** *A(LDA,\*)*, *B(LDB,\*)*

## Description

The **STRMM**, **DTRMM**, **CTRMM**, or **ZTRMM** subroutine performs one of the matrix-matrix operations:

$B := \alpha * \text{op}(A) * B$

OR

$B := \alpha * B * \text{op}(A)$

where  $\alpha$  is a scalar,  $B$  is an  $M$  by  $N$  matrix,  $A$  is a unit, or non-unit, upper or lower triangular matrix, and  $\text{op}(A)$  is either  $\text{op}(A) = A$  or  $\text{op}(A) = A^T$ .

## Parameters

*SIDE* On entry, *SIDE* specifies whether  $\text{op}(A)$  multiplies  $B$  from the left or right as follows:

*SIDE* = 'L' or 'l'

$B := \alpha * \text{op}(A) * B$

*SIDE* = 'R' or 'r'

$B := \alpha * B * \text{op}(A)$

Unchanged on exit.

*UPLO* On entry, *UPLO* specifies whether the matrix  $A$  is an upper or lower triangular matrix as follows:

*UPLO* = 'U' or 'u'

$A$  is an upper triangular matrix.

*UPLO* = 'L' or 'l'

$A$  is a lower triangular matrix.

Unchanged on exit.

*TRANSA* On entry, *TRANSA* specifies the form of  $\text{op}(A)$  to be used in the matrix multiplication as follows:

*TRANSA* = 'N' or 'n'

$\text{op}(A) = A$

*TRANSA* = 'T' or 't'

$\text{op}(A) = A^T$

*TRANSA* = 'C' or 'c'

$\text{op}(A) = A^T$

Unchanged on exit.

*DIAG* On entry, *DIAG* specifies whether or not  $A$  is unit triangular as follows:

*DIAG* = 'U' or 'u'

$A$  is assumed to be unit triangular.

*DIAG* = 'N' or 'n'

$A$  is not assumed to be unit triangular.

Unchanged on exit.

*M* On entry, *M* specifies the number of rows of  $B$ ; *M* must be at least 0; unchanged on exit.

*N* On entry, *N* specifies the number of columns of  $B$ ; *N* must be at least 0; unchanged on exit.

*ALPHA* On entry, *ALPHA* specifies the scalar  $\alpha$ . When  $\alpha$  is 0 then  $A$  is not referenced and  $B$  need not be set before entry; unchanged on exit.

<i>A</i>	An array of dimension ( <i>LDA</i> , <i>k</i> ), where <i>k</i> is <i>M</i> when <i>SIDE</i> = 'L' or 'l' and is <i>N</i> when <i>SIDE</i> = 'R' or 'r'; on entry with <i>UPLO</i> = 'U' or 'u', the leading <i>k</i> by <i>k</i> upper triangular part of the array <i>A</i> must contain the upper triangular matrix and the strictly lower triangular part of <i>A</i> is not referenced; on entry with <i>UPLO</i> = 'L' or 'l', the leading <i>k</i> by <i>k</i> lower triangular part of the array <i>A</i> must contain the lower triangular matrix and the strictly upper triangular part of <i>A</i> is not referenced. When <i>DIAG</i> = 'U' or 'u', the diagonal elements of <i>A</i> are not referenced either, but are assumed to be unity; unchanged on exit.
<i>LDA</i>	On entry, <i>LDA</i> specifies the first dimension of <i>A</i> as declared in the calling (sub) program. When <i>SIDE</i> = 'L' or 'l' then <i>LDA</i> must be at least max( 1, <i>M</i> ), when <i>SIDE</i> = 'R' or 'r' then <i>LDA</i> must be at least max( 1, <i>N</i> ); unchanged on exit.
<i>B</i>	An array of dimension ( <i>LDB</i> , <i>N</i> ); on entry, the leading <i>M</i> by <i>N</i> part of the array <i>B</i> must contain the matrix <i>B</i> , and on exit is overwritten by the transformed matrix.
<i>LDB</i>	On entry, <i>LDB</i> specifies the first dimension of <i>B</i> as declared in the calling (sub) program; <i>LDB</i> must be at least max( 1, <i>M</i> ); unchanged on exit.

## STRSM, DTRSM, CTRSM, or ZTRSM Subroutine

### Purpose

Solves certain matrix equations.

### Library

BLAS Library (**libblas.a**)

### FORTRAN Syntax

```
SUBROUTINE STRSM(SIDE, UPLO, TRANSA, DIAG,
M, N, ALPHA, A, LDA, B, LDB)
```

```
CHARACTER*1 SIDE, UPLO, TRANSA, DIAG
```

```
INTEGER M, N, LDA, LDB
```

```
REAL ALPHA
```

```
REAL A(LDA,*), B(LDB,*)
```

```
SUBROUTINE DTRSM(SIDE, UPLO, TRANSA, DIAG,
M, N, ALPHA, A, LDA, B, LDB)
```

```
CHARACTER*1 SIDE, UPLO, TRANSA, DIAG
```

```
INTEGER M, N, LDA, LDB
```

```
DOUBLE PRECISION ALPHA
```

```
DOUBLE PRECISION A(LDA,*), B(LDB,*)
```

```
SUBROUTINE CTRSM(SIDE, UPLO, TRANSA, DIAG,
M, N, ALPHA, A, LDA, B, LDB)
```

```
CHARACTER*1 SIDE, UPLO, TRANSA, DIAG
```

```
INTEGER M, N, LDA, LDB
```

```
COMPLEX ALPHA
```

```
COMPLEX A(LDA,*), B(LDB,*)
```

```
SUBROUTINE ZTRSM(SIDE, UPLO, TRANSA, DIAG,
M, N, ALPHA, A, LDA, B, LDB)
```

```
CHARACTER*1 SIDE, UPLO, TRANSA, DIAG
```

```
INTEGER M, N, LDA, LDB
```

```
COMPLEX*16 ALPHA
```

```
COMPLEX*16 A(LDA,*), B(LDB,*)
```

### Description

The **STRSM**, **DTRSM**, **CTRSM**, or **ZTRSM** subroutine solves one of the matrix equations:

- $op(A) * X = alpha * B$
- $X * op(A) = alpha * B$

where alpha is a scalar, X and B are M by N matrices, A is a unit, or non-unit, upper or lower triangular matrix, and op( A ) is either op( A ) = A or op( A ) = A'. The matrix X is overwritten on B.

## Parameters

*SIDE* On entry, *SIDE* specifies whether op( A ) appears on the left or right of X as follows:

*SIDE* = 'L' or 'l'  
 op( A ) \* X = alpha \* B

*SIDE* = 'R' or 'r'  
 X \* op( A ) = alpha \* B

Unchanged on exit.

*UPLO* On entry, *UPLO* specifies whether the matrix A is an upper or lower triangular matrix as follows:

*UPLO* = 'U' or 'u'  
 A is an upper triangular matrix.

*UPLO* = 'L' or 'l'  
 A is a lower triangular matrix.

Unchanged on exit.

*TRANS* On entry, *TRANS* specifies the form of op( A ) to be used in the matrix multiplication as follows:

*TRANS* = 'N' or 'n'  
 op( A ) = A

*TRANS* = 'T' or 't'  
 op( A ) = A'

*TRANS* = 'C' or 'c'  
 op( A ) = A'

Unchanged on exit.

*DIAG* On entry, *DIAG* specifies whether or not A is unit triangular as follows:

*DIAG* = 'U' or 'u'  
 A is assumed to be unit triangular.

*DIAG* = 'N' or 'n'  
 A is not assumed to be unit triangular.

Unchanged on exit.

*M* On entry, *M* specifies the number of rows of B; *M* must be at least 0; unchanged on exit.

*N* On entry, *N* specifies the number of columns of B; *N* must be at least 0; unchanged on exit.

*ALPHA* On entry, *ALPHA* specifies the scalar alpha. When alpha is 0 then A is not referenced and B need not be set before entry; unchanged on exit.

*A* An array of dimension ( LDA, k ), where k is M when *SIDE* = 'L' or 'l' and is N when *SIDE* = 'R' or 'r'. On entry with *UPLO* = 'U' or 'u', the leading k by k upper triangular part of the array A must contain the upper triangular matrix and the strictly lower triangular part of A is not referenced; on entry with *UPLO* = 'L' or 'l', the leading k by k lower triangular part of the array A must contain the lower triangular matrix and the strictly upper triangular part of A is not referenced. When *DIAG* = 'U' or 'u', the diagonal elements of A are not referenced, but are assumed to be unity; unchanged on exit.

*LDA* On entry, *LDA* specifies the first dimension of A as declared in the calling (sub) program. When *SIDE* = 'L' or 'l', *LDA* must be at least max( 1, M ); when *SIDE* = 'R' or 'r', *LDA* must be at least max( 1, N ); unchanged on exit.

*B* An array of dimension ( LDB, N ); on entry, the leading M by N part of the array B must contain the right-hand side matrix B, and on exit is overwritten by the solution matrix X.

*LDB* On entry, *LDB* specifies the first dimension of B as declared in the calling (sub) program. *LDB* must be at least max( 1, M ); unchanged on exit.



---

## Appendix A. Base Operating System Error Codes for Services That Require Path-Name Resolution

The following errors apply to any service that requires path name resolution:

<b>EACCES</b>	Search permission is denied on a component of the path prefix.
<b>EFAULT</b>	The <i>Path</i> parameter points outside of the allocated address space of the process.
<b>EIO</b>	An I/O error occurred during the operation.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the <i>Path</i> parameter.
<b>ENAMETOOLONG</b>	A component of a path name exceeded 255 characters and the process has the <b>DisallowTruncation</b> attribute (see the <b>ulimit</b> subroutine) or an entire path name exceeded 1023 characters.
<b>ENOENT</b>	A component of the path prefix does not exist.
<b>ENOENT</b>	A symbolic link was named, but the file to which it refers does not exist.
<b>ENOENT</b>	The path name is null.
<b>ENOTDIR</b>	A component of the path prefix is not a directory.
<b>ESTALE</b>	The root or current directory of the process is located in a virtual file system that is unmounted.

### Related Information

List of File and Directory Manipulation Services.



---

## Appendix B. ODM Error Codes

When an ODM subroutine is unsuccessful, a value of -1 is returned and the **odmerrno** variable is set to one of the following values:

<b>ODMI_BAD_CLASSNAME</b>	The specified object class name does not match the object class name in the file. Check path name and permissions.
<b>ODMI_BAD_CLXNNAME</b>	The specified collection name does not match the collection name in the file.
<b>ODMI_BAD_CRIT</b>	The specified search criteria is incorrectly formed. Make sure the criteria contains only valid descriptor names and the search values are correct. For information on qualifying criteria, see "Understanding ODM Object Searches" in <i>AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs</i> .
<b>ODMI_BAD_LOCK</b>	Cannot set a lock on the file. Check path name and permissions.
<b>ODMI_BAD_TIMEOUT</b>	The time-out value was not valid. It must be a positive integer.
<b>ODMI_BAD_TOKEN</b>	Cannot create or open the lock file. Check path name and permissions.
<b>ODMI_CLASS_DNE</b>	The specified object class does not exist. Check path name and permissions.
<b>ODMI_CLASS_EXISTS</b>	The specified object class already exists. An object class must not exist when it is created.
<b>ODMI_CLASS_PERMS</b>	The object class cannot be opened because of the file permissions.
<b>ODMI_CLXNMAGICNO_ERR</b>	The specified collection is not a valid object class collection.
<b>ODMI_FORK</b>	Cannot fork the child process. Make sure the child process is executable and try again.
<b>ODMI_INTERNAL_ERR</b>	An internal consistency problem occurred. Make sure the object class is valid or contact the person responsible for the system.
<b>ODMI_INVALID_CLASS</b>	The specified file is not an object class.
<b>ODMI_INVALID_CLXN</b>	Either the specified collection is not a valid object class collection or the collection does not contain consistent data.
<b>ODMI_INVALID_PATH</b>	The specified path does not exist on the file system. Make sure the path is accessible.
<b>ODMI_LINK_NOT_FOUND</b>	The object class that is accessed could not be opened. Make sure the linked object class is accessible.
<b>ODMI_LOCK_BLOCKED</b>	Cannot grant the lock. Another process already has the lock.
<b>ODMI_LOCK_ENV</b>	Cannot retrieve or set the lock environment variable. Remove some environment variables and try again.
<b>ODMI_LOCK_ID</b>	The lock identifier does not refer to a valid lock. The lock identifier must be the same as what was returned from the <b>odm_lock</b> subroutine.
<b>ODMI_MAGICNO_ERR</b>	The class symbol does not identify a valid object class.
<b>ODMI_MALLOC_ERR</b>	Cannot allocate sufficient storage. Try again later or contact the person responsible for the system.
<b>ODMI_NO_OBJECT</b>	The specified object identifier did not refer to a valid object.
<b>ODMI_OPEN_ERR</b>	Cannot open the object class. Check path name and permissions.
<b>ODMI_OPEN_PIPE</b>	Cannot open a pipe to a child process. Make sure the child process is executable and try again.
<b>ODMI_PARAMS</b>	The parameters passed to the subroutine were not correct. Make sure there are the correct number of parameters and that they are valid.
<b>ODMI_READ_ONLY</b>	The specified object class is opened as read-only and cannot be modified.
<b>ODMI_READ_PIPE</b>	Cannot read from the pipe of the child process. Make sure the child process is executable and try again.
<b>ODMI_TOOMANYCLASSES</b>	Too many object classes have been accessed. An application can only access less than 1024 object classes.
<b>ODMI_UNLINKCLASS_ERR</b>	Cannot remove the object class from the file system. Check path name and permissions.
<b>ODMI_UNLINKCLXN_ERR</b>	Cannot remove the object class collection from the file system. Check path name and permissions.
<b>ODMI_UNLOCK</b>	Cannot unlock the lock file. Make sure the lock file exists.

## **Related Information**

List of ODM Commands and Subroutines in *AIX 5L Version 5.2 General Programming Concepts: Writing and Debugging Programs*.

---

## Appendix C. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
Dept. LRAS/Bldg. 003  
11400 Burnet Road  
Austin, TX 78758-3498  
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106, Japan

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

---

## Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be the trademarks or service marks of others.

---

# Index

## Special characters

`_lazySetErrorHandler` subroutine 584  
`_showstring` subroutine 626  
`_sync_cache_range` subroutine 311

## Numerics

8-bit character capability 588

## A

absolute values  
  finding index of element with maximum value 669  
access control information  
  retrieving 272  
access control subroutines  
  fstatacl 272  
  revoke 45  
  statacl 272  
addch subroutine 523  
addstr subroutine 524  
alarm signals  
  beeping 530  
  flashing 557  
alphasort subroutine 107  
alternate stack 174  
argument formatting  
  vfprintf 432  
  vscanf 432  
  vsscanf 432  
asynchronous serial data line  
  sending breaks on 348  
atoi subroutine 299  
atrtroff subroutine 526  
attron subroutine 528  
attrset subroutine 528  
authentication database  
  opening and closing 148  
authentication subroutines  
  endpwnb 148  
  enduserdb 153  
  setpwnb 148  
  setuserdb 153  
  tcb 341

## B

backspace character  
  returning 556  
baudrate subroutine 529  
beep subroutine 530  
Berkeley Compatibility Library  
  subroutines  
    rand\_r 6  
binary trees, manipulating 410  
BLAS matrix-matrix operations 711

BLAS matrix-matrix subroutines 705, 706, 708, 710,  
  713, 714, 716, 718  
BLAS matrix-vector subroutines 673, 674, 676, 677,  
  679, 680, 681, 682, 683, 685, 687, 688, 690, 692,  
  694, 695, 696, 697, 698, 699, 700, 701, 702, 704  
BLAS vector-vector functions 661, 662, 667, 668, 669,  
  670  
BLAS vector-vector subroutines 663, 664, 665, 666,  
  668, 671, 672  
box subroutine 531  
buffers  
  assigning to streams 129  
bytes  
  copying 304

## C

carriage return 598  
CAXPY subroutine 663  
cbox subroutine 531  
cboxalt subroutine 531  
CBREAK mode 535  
cbreak subroutine 535  
CCOPY subroutine 665  
CDOTC function 661  
CDOTU function 662  
CGBMV subroutine 674  
CGEMM subroutine 705  
CGEMV subroutine 673  
CGERC subroutine 695  
CGERU subroutine 695  
change color definition 575  
change color-pair definition 576  
change terminal capabilities 557  
character conversion  
  wide characters  
    lowercase to uppercase 367  
    to double-precision number 455  
    to long integer 459  
    to multibyte 461, 468  
    to tokens 458  
    to unsigned long integer 463  
    uppercase to lowercase 366  
character data  
  interpreting 109  
  reading 109  
character manipulation subroutines  
  vwsprintf 439  
character mapping 469  
character transliteration 365  
characters  
  adding  
    lines 579  
    single characters 523, 578  
    strings 524  
  backspace 556  
  clearing screen 536, 537  
  controlling text scrolling 619, 620, 622

- characters (*continued*)
  - deleting 550
  - dumping strings 626
  - echoing 552
  - erasing lines 540, 541, 551
  - erasing window 555
  - getting single characters 561
  - getting strings 566
  - handling input 588, 599
  - line-kill 583
  - placing at cursor location 574
  - reading formatted input 614
  - refreshing 648, 650
  - type ahead 653
  - typeahead 558
  - writing 439
  - writing formatted output 604
- charsetID
  - wide character 448
- CHBMV subroutine 677
- CHEMM subroutine 708
- CHEMV subroutine 676
- CHER subroutine 696
- CHER2 subroutine 698
- CHER2K subroutine 714
- CHERK subroutine 711
- CHPMV subroutine 679
- CHPR subroutine 697
- CHPR2 subroutine 699
- clear subroutine 536
- clearok subroutine 537
- close role database 149
- close SMIT ACL database 127
- closelog subroutine 329
- closelog\_r subroutine 332
- clrtobot subroutine 540
- clrtoeol subroutine 541
- code sets
  - reading map files 131
- color definition 575
- color intensity 542
- color manipulation 532
- color pair 602
- color support 570
- color-pair definition 576
- color, initialize 636
- columns
  - determining number 624, 640
- compare wide character 501
- complex dot products
  - determining 661, 662
- control characters
  - specifying 654
- control input characters 570
- convert wide character 468
- converter subroutines
  - wcsrtombs 452
- copy a window region 543
- copy wide character 501, 502
- create subwindows 637
- cresetty subroutine 611
- CROTG subroutine 663
- CSCAL subroutine 668
- CSROT subroutine 664
- CSSCAL subroutine 668
- CSWAP subroutine 666
- CSYMM subroutine 706
- CSYR2K subroutine 713
- CSYRK subroutine 710
- CTBMV subroutine 685
- CTBSV subroutine 690
- CTPMV subroutine 687
- CTPSV subroutine 692
- CTRMM subroutine 716
- CTRMV subroutine 683
- CTRSM subroutine 718
- CTRSV subroutine 688
- current process credentials
  - setting 139
- current process environment
  - setting 142
- current processes
  - group ID
    - setting 134
  - suspending 196
  - user information 426
- current screen
  - refreshing 608, 657
- current screens
  - refreshing 603
- curses
  - initializing 577
  - terminating 554
- curses character control subroutines
  - \_showstring 626
  - addch 523
  - addstr 524
  - clear 536
  - clearok 537
  - clrtobot 540
  - clrtoeol 541
  - delch 550
  - deleteln 551
  - erase 555
  - getch 561
  - getstr 566
  - inch 574
  - insch 578
  - insertln 579
  - meta 588
  - mvaddch 523
  - mvaddstr 524
  - mvdelch 550
  - mvgetch 561
  - mvgetstr 566
  - mvinch 574
  - mvinsch 578
  - mvscanw 614
  - mvwaddch 523
  - mvwaddstr 524
  - mvwdelch 550
  - mvwgetch 561

curses character control subroutines (*continued*)

- mvwgetstr 566
- mvwinch 574
- mvwisch 578
- mvwscanw 614
- nodelay 599
- scanw 614
- scroll 619
- scrollok 620
- setscreg 622
- unctrl 654
- waddch 523
- waddstr 524
- wclear 536
- wclrtoebot 540
- wclrtoeol 541
- wdeich 550
- wdeleteln 551
- werase 555
- wgetch 561
- wgetstr 566
- winch 574
- wisch 578
- winsertln 579
- wscanw 614
- wsetscreg 622

curses cursor control subroutines

- getyx 569
- leaveok 585
- move 589
- mvcur 589
- wmove 589

curses data structure 617

curses options setting subroutines

- idlok 573
- intrlush 580
- keypad 582
- typeahead 653

curses portability subroutines

- baudrate 529
- erasechar 556
- flushinp 558
- killchar 583

curses subroutine

- getbegyx 560
- getmaxyx 565

curses subroutines

- character locations
  - echochar, wechochar, pechochar 553
- endwin 554
- initscr 577
- switching input/output to different terminals 624

curses terminal manipulation subroutines

- cbreak 535
- crsetty 611
- def\_prog\_mode 545
- def\_shell\_mode 546
- delay\_output 549
- echo 552
- has\_ic 571
- has\_il 572

curses terminal manipulation subroutines (*continued*)

- longname 586
- newterm 594
- nl 598
- nocbreak 535
- noecho 552
- nonl 598
- noraw 607
- putp 606
- raw 607
- reset\_prog\_mode 609
- reset\_shell\_mode 610
- resetterm 610
- resetty 611
- set\_term 624
- setupterm 624
- tgetent 640
- tgetflag 641
- tgetnum 642
- tgetstr 642
- tgoto 643
- tparam 651
- tputs 652

curses video attributes subroutines

- attroff 526
- attron 528
- attrset 528
- beep 530
- flash 557
- standend 635
- standout 635
- vidattr 656
- vidputs 656
- wattroff 526
- wattron 528
- wattrset 528
- wstandend 635
- wstandout 635

curses window manipulation subroutines

- box 531
- delwin 552
- doupdate 657
- makenew 587
- mvwin 591
- newpad 592
- newwin 596
- overlay 601
- overwrite 601
- pnoutrefresh 603
- prefresh 603
- refresh 608
- subwin 638
- touchline 648
- touchoverlap 649
- touchwin 650
- wnoutrefresh 657
- wrefresh 608

cursor control

- moving logical cursor 589
- moving physical cursor 589
- placing cursor 585

cursor control (*continued*)  
    returning logical cursor coordinates 569  
cursor coordinates 560  
cursor visibility 544

## D

D cache 311  
DASUM subroutine 668  
data  
    sorting with quicker-sort algorithms 1  
data sorting subroutines  
    qsort 1  
    tdelete 410  
    tfind 410  
    tsearch 410  
    twalk 410  
data transmissions  
    suspending 343  
    waiting for completion 342  
data words  
    trace 400  
databases  
    authentication  
        opening and closing 148  
date  
    format conversions 288  
date format conversions 301, 447  
DAXPY subroutine 663  
DCOPY subroutine 665  
DDOT function 661  
def\_prog\_mode subroutine 545  
def\_shell\_mode subroutine 546  
defect 220643 364  
define character mapping 469  
delay mode 570  
delay\_output subroutine 549  
delch subroutine 550  
deleteln subroutine 551  
delwin subroutine 552  
determine terminal color support 570  
device driver  
    calling 316  
device switch tables  
    checking entry status 325  
DGBMV subroutine 674  
DGEMM subroutine 705  
DGEMV subroutine 673  
DGER subroutine 694  
directories  
    reading 21  
    removing 48  
    removing entries 423  
    renaming 42  
    scanning contents 107  
    sorting contents 107  
directory subroutines  
    alphasort 107  
    readlink 22  
    rmdir 48  
    scandir 107

directory subroutines (*continued*)  
    symlink 308  
    unlink 423  
disable terminal capabilities 557  
discard lines in windows 559  
disk quotas  
    manipulating 2  
DNRM2 function 667  
dot products  
    determining 661, 670  
doupdate subroutine 657  
drawbox subroutine 531  
drawboxalt subroutine 531  
DROT subroutine 664  
DROTG subroutine 663  
DROTM subroutine 671  
DROTMG subroutine 672  
DSBMV subroutine 681  
DSCAL subroutine 668  
DSPMV subroutine 682  
DSPR subroutine 701  
DSPR2 subroutine 704  
DSWAP subroutine 666  
DSYMM subroutine 706  
DSYMV subroutine 680  
DSYR subroutine 700  
DSYR2 subroutine 702  
DSYR2K subroutine 713  
DSYRK subroutine 710  
DTBMV subroutine 685  
DTBSV subroutine 690  
DTPMV subroutine 687  
DTPSV subroutine 692  
DTRMM subroutine 716  
DTRMV subroutine 683  
DTRSM subroutine 718  
DTRSV subroutine 688  
dump file, data structure 617  
dump file, restore screen 618  
DZASUM subroutine 668  
DZNRM2 function 667

## E

echo subroutine 552  
echochar subroutine 553  
echoing characters 552  
endpwdb subroutine 148  
endroledb subroutine 149  
enduserdb subroutine 153  
endwin subroutine 554  
equations  
    solving systems 688, 690, 692  
erase subroutine 555  
erasechar subroutine 556  
error codes 721  
error codes, ODM 723  
error handler, install 584  
error handling  
    controlling system logs 329  
    numbering error message string 285

- errorlogging subroutines
  - closelog 329
  - openlog 329
  - setlogmask 329
  - syslog 329
- errorlogging\_r subroutines 332
- Euclidean lengths
  - determining 667
- examine state of alternate stack 174
- execution control
  - saving and restoring context 135
- execution control subroutines
  - longjmp 135
  - setjmp 135
- exponential numbers
  - scalbln 106
  - scalblnf 106
  - scalblnl 106
  - scalbn 106
  - scalbnf 106
  - scalbnl 106
- extended curses
  - initializing 577
- extended curses character control subroutines
  - \_showstring 626
  - getch 561
  - inch 574
  - insch 578
  - meta 588
  - mvgetch 561
  - mvinch 574
  - mvinsch 578
  - mvscanw 614
  - mvwgetch 561
  - mvwinch 574
  - mvwinsch 578
  - mvwscanw 614
  - printw 604
  - scanw 614
  - scroll 619
  - scrollok 620
  - wgetch 561
  - winch 574
  - winsch 578
  - wscanw 614
- extended curses options setting subroutines
  - idlok 573
  - intrflush 580
- extended curses portability subroutines
  - baudrate 529
  - erasechar 556
  - flushinp 558
  - killchar 583
- extended curses subroutines
  - initscr 577
- extended curses terminal manipulation subroutines
  - delay\_output 549
  - has\_ic 571
  - has\_il 572
  - newterm 594
  - putp 606
- extended curses terminal manipulation subroutines
  - (continued)
  - resetterm 610
  - set\_term 624
  - setupterm 624
  - tgentent 640
  - tgetflag 641
  - tgetnum 642
  - tparm 651
- extended curses video attributes subroutines
  - attroff 526
  - attron 528
  - attrset 528
  - standend 635
  - standout 635
  - vidputs 656
  - wattroff 526
  - wattron 528
  - wattrset 528
  - wstandend 635
  - wstandout 635
- extended curses window manipulation subroutines
  - box 531
  - cbox 531
  - cboxalt 531
  - delwin 552
  - doupdate 657
  - drawbox 531
  - drawboxalt 531
  - fullbox 531
  - makenew 587
  - mvwin 591
  - newwin 596
  - overlay 601
  - overwrite 601
  - superbox 531
  - superbox1 531
  - touchline 648
  - touchoverlap 650
  - wnoutrefresh 657

## F

- ffullstat subroutine 277
- file access times
  - setting 428
- file creation masks
  - getting or setting values 418
- file descriptors
  - checking I/O status 115
- file modification times
  - setting 428
- file subroutines
  - ffullstat 277
  - fstat 277
  - fstatx 277
  - fttruncate 407
  - fullstat 277
  - lstat 277
  - remove 40
  - rename 42

- file subroutines (*continued*)
  - stat 277
  - statx 277
  - tempnam 363
  - tmpfile 363
  - tmpnam 363
  - truncate 407
  - umask 418
  - utime 428
  - utimes 428
- file system information 276
- file system subroutines
  - fstatfs 274
  - mount 435
  - quotactl 2
  - statfs 274
  - sync 310
  - sysconf 311
  - umount 419
  - ustat 274
  - uvmount 419
  - vmount 435
- file systems
  - manipulating disk quotas 2
  - mounting 435
  - returning statistics 274
  - unmounting 419
  - updating 310
- file, input/output 615
- files
  - changing length of regular 407
  - constructing names for temporary 363
  - creating symbolic links 308
  - creating temporary 363
  - deleting 40
  - providing status information 277
  - reading 16
  - removing 40
  - renaming 42
  - revoking access 45
  - writing to 506
- find wide character 500
- find wide character substring 454
- flash subroutine 557
- flow control
  - performing 343
- flushing
  - typeahead characters 558
- flushinp subroutine 558
- foreground process group IDs
  - getting 347
  - setting 351
- formatted input
  - converting 109
- fscanf subroutine 109
- fstat subroutine 277
- fstatacl subroutine 272
- fstatfs subroutine 274
- fstatvfs subroutine 276
- fstatx subroutine
  - described 277
- ftruncate subroutine 407
- fullbox subroutine 531
- fullstat subroutine 277

## G

- gamma subroutines
  - tgamma 354
  - tgammaf 354
  - tgammaL 354
- get capabilities, terminfo 644
- get key name 581
- get terminals numeric value 646
- get terminals string capability 647
- get XTI variables 379
- get\_wctype subroutine 470
- getbegyx subroutine 560
- getch subroutine 561, 588, 599
- getmaxyx subroutine 565
- getstr subroutine 566
- getyx macro 569
- Givens plane rotations
  - constructing 663
- Givens transformations
  - applying 671
  - constructing 672
- gsignal subroutine 271
- gtty subroutine 303

## H

- half-delay mode 570
- has\_ic subroutine 571
- has\_il subroutine 572
- Hermitian operations
  - performing rank 1 696, 697
  - performing rank 2 698, 699
  - performing rank 2k 714
  - performing rank k 711
- highlight mode 635
- hook words
  - trace 400
- hyperbolic functions
  - computing 194
- hyperbolic sine subroutines
  - sinhf 194
- hyperbolic tangent subroutines
  - tanhf 340

## I

- I cache 311
- I/O asynchronous subroutines
  - select 115
- I/O low-level subroutines 16, 506
  - readvx 16
  - readx 16
  - writevx 506
  - writex 506
- I/O stream subroutines
  - fscanf 109

## I/O stream subroutines (*continued*)

- scanf 109
- setbuf 129
- setbuffer 129
- setlinebuf 129
- setvbuf 129
- sscanf 109
- ungetc 422
- ungetwc 422
- wscanf 109

## I/O terminal subroutines

- gtty 303
- isatty 413
- stty 303
- tcdrain 342
- tcflow 343
- tcflush 344
- tcgetattr 346
- tcgetpgrp 347
- tcsendbreak 348
- tcsetattr 349
- tcsetpgrp 351
- termdef 352
- ttylock 412
- ttylocked 412
- ttyname 413
- ttyslot 415
- ttyunlock 412
- ttywait 412

ICAMAX subroutine 669

IDAMAX subroutine 669

idlok subroutine 573

idxpg4 281

inch subroutine 574

index subroutine 291

initialize color 636

initscr subroutine 577

initstate subroutine 7

input streams

- pushing single character into 422

insch subroutine 578

insert-character capability 571

insert-line capability 572

insert/delete line option 573

insertln subroutine 579

interval timers

- releasing 39

intrflush subroutine 580

ISAMAX subroutine 669

isatty subroutine 413

IZAMAX subroutine 669

## J

JFS

- manipulating disk quotas 2

## K

kernel configurations

- customizing 315

kernel extension modules

- loading 328

kernel extensions

- loading 321

kernel object files

- determining status 326

- invoking 318

- unloading 323

kernel parameters

- setting 327

key name 581

keypad

- enabling 582

keypad subroutine 582

killchar subroutine 583

## L

label name, return 631

lazy loading runtime system 584

LC\_ALL environment variable 137

LC\_COLLATE category 137

LC\_CTYPE category 137

LC\_MESSAGES category 137

LC\_MONETARY category 137

LC\_NUMERIC category 137

LC\_TIME category 137

leaveok subroutine 585

line-kill character 583

lines

- adding 579

- determining number 624, 640

- erasing 540, 541, 551

links

- creating symbolic 308

- reading contents of symbolic 22

locale subroutines

- rpmatch 49

- setlocale 136

locales

- changing or querying 136

- response matching 49

localization subroutines

- strfmon 286

- strftime 288

- strptime 301

locking functions

- controlling tty 412

logical cursor 569, 589

long integers, converting

- from character strings 299

- from wide-character strings 459

long numeric data 154

longjmp subroutine 135

longname subroutine 586

lowercase characters

- converting from uppercase 366

- converting to uppercase 367

lstat subroutine 277

## M

- m\_initscr subroutine 577
- makenew subroutine 587
- mapped files
  - attaching to process 155
- mapping, character 469
- matrices
  - performing matrix-matrix operations with
    - general matrices 705
    - Hermitian matrices 708
    - symmetric matrices 706
    - triangular matrices 716
  - performing matrix-vector operations with
    - general banded matrices 674
    - general matrices 673
    - Hermitian band matrices 677
    - Hermitian matrices 676
    - packed Hermitian matrices 679
    - packed symmetric matrices 682
    - packed triangular matrices 687
    - symmetric band matrices 681
    - symmetric matrices 680
    - triangular band matrices 685
    - triangular matrices 683
  - solving equations 718
- memory
  - freeing 506
- memory management
  - activating paging or swapping 305, 306
  - controlling shared memory operations 158
  - returning paging device status 307
  - returning shared memory segments 161
- memory management subroutines
  - shmat 155
  - shmctl 158
  - shmdt 160
  - shmget 161
  - swapoff 305
  - swapon 306
  - swapqry 307
- memory mapping
  - attaching segment or file to process 155
- message queues
  - checking I/O status 115
- meta subroutine 588
- minicurses
  - initializing 577
- minicurses subroutines
  - attrset 528
  - baudrate 529
  - erasechar 556
  - flushinp 558
  - getch 561
  - m\_initscr 577
- monetary strings 286
- mount subroutine 435
- mounted file systems
  - returning statistics 274
- move subroutine 589
- multibyte characters
  - converting from wide 461, 468

- mvaddch subroutine 523
- mvaddstr subroutine 524
- mvcur subroutine 589
- mvdelch subroutine 550
- mvgetch subroutine 561
- mvgetstr subroutine 566
- mvinch subroutine 574
- mvinsch subroutine 578
- mvprintw subroutine 604
- mvscanw subroutine 614
- mwaddch subroutine 523
- mwaddstr subroutine 524
- mwdelch subroutine 550
- mwvgetch subroutine 561
- mwvgetstr subroutine 566
- mvwin subroutine 591
- mvwinch subroutine 574
- mvwinsch subroutine 578
- mvwprintw subroutine 604
- mvwscanw subroutine 614

## N

- new-line character 598
- newpad subroutine 592
- newterm subroutine 594
- newwin subroutine 596
- nl subroutine 598
- no timeout mode 599
- nocbreak subroutine 535
- nodelay subroutine 599
- noecho subroutine 552
- nonl subroutine 598
- noraw subroutine 607
- nsleep subroutine 196
- numbers
  - generating
    - pseudo-random 5
    - random 5, 7
- numerical data
  - generating pseudo-random numbers 6
- numerical manipulation subroutines
  - atoi 299
  - initstate 7
  - rand 5
  - random 7
  - rnl 46
  - rint 46
  - rsqrt 104
  - scalb 106
  - setstate 7
  - sgetl 154
  - sinh 194
  - sinl 193
  - sputl 154
  - sqrt 238
  - sqrtl 238
  - srand 5
  - srandom 7
  - strtod 295
  - strtof 295

numerical manipulation subroutines (*continued*)

- strtol 299
- strtold 295
- strtoul 299
- tan 339
- tanh 340
- tanl 339
- watof 514
- watoi 515
- watol 515
- wstrtod 514
- wstrtol 515

## O

object file access subroutines

- sgetl 154
- sputl 154

object file subroutines

- unload 424

object files

- unloading 424

Obtaining high-resolution elapsed time

- read\_real\_time or time\_base\_to\_time 23

ODM error codes 723

open role database 149

open SMIT ACL database 127

openlog subroutine 329

openlog\_r subroutine 332

operating system

- customizing configurations 315
- identifying 420

output

- waiting for completion 342

overlay subroutine 601

overwrite subroutine 601

## P

paging memory

- activating 305, 306
- returning information on devices 307

parameter lists

- handling variable-length 429

parameter structures

- copying into buffers 319, 320

path name

- resolve 25

path-name resolution 721

pechochar subroutine 553

performance data from remote kernels 105

physical cursor 589

plane rotations

- applying 664

pnoutrefresh subroutine 603

prefresh subroutine 603

print formatted output 439

printf subroutine 604

printw subroutine 604

process credentials

- setting 139

process environments

- setting 142

process group IDs

- returning 347
- setting 132, 146, 150, 351
- supplementary IDs
- setting 134

process identification

- current operating system name 420

process initiation

- restarting system 26

process priorities

- setting scheduled priorities 147
- yielding to higher priorities 520

process resource allocation

- setting and getting user limits 415

process signals

- blocked signal sets

- changing 188
- returning 179

- changing subroutine restart behavior 177

- enhancement and management 183

- handling system-defined exceptions 164

- implementing software signal facility 271

- manipulating signal sets 175

- sending to executing program 4

- signal masks

- replacing 188
- saving or restoring 186
- setting 179

- specifying action upon delivery 164

- stacks

- defining alternate 187

- saving or restoring context 186

process subroutines (security and auditing)

- setegid 132
- seteuid 151
- setgid 132
- setgidx 132
- setgroups 134
- setpcred 139
- setpenv 142
- setregid 132
- setreuid 151
- setrgid 132
- setruid 151
- setuid 151
- setuidx 151
- system 337
- usrinfo 426

process user IDs

- setting 151

processes

- handling user information 426
- suspending 196, 440, 443

processes subroutines

- gsignal 271
- raise 4
- reboot 26
- semctl 120
- semget 123

processes subroutines (*continued*)

- semop 125
- setpgid 146
- setpgrp 146
- setpri 147
- setsid 150
- sigaddset 175
- sigblock 179
- sigdelset 175
- sigemptyset 175
- sigfillset 175
- sighold 183
- sigignore 183
- siginterrupt 177
- sigismember 175
- siglongjmp 186
- sigpause 188
- sigpending 179
- sigprocmask 179
- sigreize 183
- sigset 183
- sigsetjmp 186
- sigsetmask 179
- sigstack 187
- sigsuspend 188
- ssignal 271
- ulimit 415
- uname 420
- unamex 420
- wait 440
- wait3 440
- waitid 443
- waitpid 440
- yield 520

program mode 609

pseudo-random numbers

- generating 5

pthread\_kill subroutine 4

push character to input queue 655

putp subroutine 606

## Q

qsort subroutine 1

queues

- discarding data 344

quotactl subroutine 2

## R

ra\_attachrset Subroutine 9

ra\_detachrset Subroutine 11

ra\_exec Subroutine 12

ra\_fork Subroutine 14

ra\_getrset Subroutine 15

raise subroutine 4

rand subroutine 5

rand\_r subroutine 6

random numbers

- generating 5, 7

random subroutine 7

rank 1 operations 694, 695

raw mode 607

raw subroutine 607

re\_comp subroutine 28

re\_exec subroutine 28

re-initializest terminal structures 612

read operations

- from a file 16

read subroutine 16

read\_real\_time Subroutine 23

readdir\_r subroutine 21

readlink subroutine 22

readv subroutine

- described 16

readvx subroutine 16

readx subroutine

- described 16

realpath subroutine 25

reboot subroutine 26

receive data unit 371

reception of data

- suspending 343

reciprocals of square roots

- computing 104

refresh subroutine 608

refreshing

- characters 648, 650
- current screen 603, 608, 657
- standard screen 657
- terminal 603, 608
- windows 649, 658

regcmp subroutine 29

regcomp subroutine 32

regerror subroutine 33

regex subroutine 29

regexec subroutine 35

regfree subroutine 38

regular expression subroutines

- regcmp 29
- regcomp 32
- regerror 33
- regex 29
- regexec 35
- regfree 38

regular expressions

- comparing 35
- compiling 29, 32
- error messages 33
- freeing memory 38
- matching 29

regular files

- changing length 407

release indication

- user data 367

relinquish processor 115

reltimerid subroutine 39

remainder subroutine 39

remainder subroutines

- remquo 41
- remquof 41
- remquol 41

Remainder subroutines  
 remainder 39  
 remainderf 39  
 remainderl 39  
 remainderf subroutine 39  
 remainderl subroutine 39  
 remote hosts  
 rstat subroutine 105  
 Remote Statistics Interface subroutines  
 RSiChangeFeed 53  
 RSiChangeHotFeed 54  
 RSiClose 55  
 RSiCreateStatSet 57  
 RSiDelSetHot 58  
 RSiDelSetStat 59  
 RSiFirstCx 60  
 RSiFirstStat 62  
 RSiGetHotItem 63  
 RSiGetRawValue 65  
 RSiGetValue 66  
 RSiInit 68  
 RSiInstantiate 69  
 RSiMainLoop 72  
 RSiNextCx 73  
 RSiNextStat 74  
 RSiOpen 75  
 RSiPathAddSetStat 78  
 RSiPathGetCx 79  
 RSiStartFeed 80  
 RSiStartHotFeed 81  
 RSiStatGetPath 82  
 RSiStopHotFeed 85  
 remove subroutine 40  
 remquo subroutine 41  
 remquof subroutine 41  
 remquol subroutine 41  
 rename subroutine 42  
 replace lines in windows 559  
 reserve a screen line 612  
 reset\_malloc\_log subroutine 44  
 reset\_prog\_mode subroutine 609  
 reset\_shell\_mode subroutine 610  
 resetterm subroutine 610  
 resetty subroutine 611  
 Resource Set APIs  
 ra\_attachrset 9  
 ra\_detachrset 11  
 ra\_exec 12  
 ra\_fork 14  
 ra\_getrset 15  
 rs\_alloc 86  
 rs\_discardname 87  
 rs\_free 88  
 rs\_getassociativity 89  
 rs\_getinfo 90  
 rs\_getnameattr 91  
 rs\_getnamedrset 92  
 rs\_getpartition 93  
 rs\_getrad 94  
 rs\_init 96

Resource Set APIs (continued)  
 rs\_numrads 96  
 rs\_op 97  
 rs\_registername 99  
 rs\_setnameattr 101  
 rs\_setpartition 103  
 restore soft function key 633  
 restore virtual screen 618  
 retrieves information from terminfo 547  
 return color intensity 542  
 return file system information 276  
 return label, soft label 631  
 return window size 565  
 returns color to color pair 602  
 revoke subroutine 45  
 rindex subroutine 291  
 rint subroutine 46  
 rintf subroutine 46  
 rintl subroutine 46  
 ripoffline subroutine 612  
 rmdir subroutine 48  
 round subroutine 47  
 roundf subroutine 47  
 rounding numbers  
 rintf 46  
 rintl 46  
 round 47  
 roundf 47  
 roundl 47  
 trunc 406  
 truncf 406  
 trunci 406  
 roundl subroutine 47  
 rpmatch subroutine 49  
 rs\_alloc Subroutine 86  
 rs\_discardname Subroutine 87  
 rs\_free Subroutine 88  
 rs\_getassociativity Subroutine 89  
 rs\_getinfo Subroutine 90  
 rs\_getnameattr Subroutine 91  
 rs\_getnamedrset Subroutine 92  
 rs\_getpartition Subroutine 93  
 rs\_getrad Subroutine 94  
 rs\_init Subroutine 96  
 rs\_numrads Subroutine 96  
 rs\_op Subroutine 97  
 rs\_registername Subroutine 99  
 rs\_setnameattr Subroutine 101  
 rs\_setpartition Subroutine 103  
 rsqrt subroutine 104  
 rstat subroutine 105  
 runtime tunable parameters  
 setting 327

## S

SASUM subroutine 668  
 savetty subroutine 614  
 SAXPY subroutine 663  
 scalb subroutine 106  
 scalbln subroutine 106

- scalblnf subroutine 106
- scalblnl subroutine 106
- scalbn subroutine 106
- scalbnf subroutine 106
- scalbnl subroutine 106
- scandir subroutine 107
- scanf subroutine 109, 614
- scanw subroutine 614
- SCASUM subroutine 668
- sched\_yield subroutine 115
- scheduling policy and priority
  - kernel thread 360
- SCNRM2 function 667
- SCOPY subroutine 665
- scr\_dump subroutine 615
- scr\_init subroutine 617
- scr\_restore subroutine 618
- screen line 612
- screens
  - refreshing 603, 608, 658
- scroll subroutine 619
- scrollok subroutine 620
- SDOT function 661
- SDSDOT function 670
- select subroutine 115
- semaphore identifiers 123
- semaphore operations 120, 125
- semctl subroutine 120
- semget subroutine 123
- semop subroutine 125
- send data 373
- serial data lines
  - sending breaks on 348
- sessions
  - creating 150
- set blocking or non-blocking read 599
- set cursor visibility 544
- set terminal variables 621
- set wide character 503
- set\_curterm subroutine 621
- set\_term subroutine 624
- setauthdb subroutine 128
- setbuf subroutine 129
- setbuffer subroutine 129
- setcsmap subroutine 131
- setegid subroutine 132
- seteuid subroutine 151
- setgid subroutine 132
- setgidx subroutine 132
- setgroups subroutine 134
- setjmp subroutine 135
- setlinebuf subroutine 129
- setlocale subroutine 136
- setlogmask subroutine 329
- setlogmask\_r subroutine 332
- setpcred subroutine 139
- setpenv subroutine 142
- setpgid subroutine 146
- setpgrp subroutine 146
- setpri subroutine 147
- setpwdb subroutine 148
- setregid subroutine 132
- setreuid subroutine 151
- setrgid subroutine 132
- setroledb subroutine 149
- setruid subroutine 151
- setscrreg subroutine 622
- setsid subroutine 150
- setstate subroutine 7
- setsyx subroutine 623
- setuid subroutine 151
- setuidx subroutine 151
- setup soft labels 633
- setupterm subroutine 624
- setuserdb subroutine 153
- setvbuf subroutine 129
- SGBMV subroutine 674
- SGEMM subroutine 705
- SGEMV subroutine 673
- SGER subroutine 694
- sgetl subroutine 154
- shared memory segments
  - attaching to process 155
  - detaching 160
  - operations on 158
  - returning 161
- shell commands
  - running 337
- shell mode 546, 610
- shmat subroutine 155
- shmctl subroutine 158
- shmdt subroutine 160
- shmget subroutine 161
- short status requests
  - sending 259, 262
- sigaddset subroutine 175
- sigaltstack subroutine 174
- sigblock subroutine 179
- sigdelset subroutine 175
- sigemptyset subroutine 175
- sigfillset subroutine 175
- sighold subroutine 183
- sigignore subroutine 183
- siginterrupt subroutine 177
- sigismember subroutine 175
- siglongjmp subroutine 186
- signal masks
  - replacing 188
  - saving or restoring 186
  - setting 179
- signal stacks
  - defining alternate 187
  - saving or restoring context 186
- signbit macro 178
- sigpause subroutine 188
- sigpending subroutine 179
- sigprocmask subroutine 179
- sigqueue subroutine 181
- sigrelse subroutine 183
- sigset subroutine 183
- sigsetjmp subroutine 186
- sigsetmask subroutine 179

sigstack subroutine 187  
 sigsuspend subroutine 188  
 sigtimedwait subroutine 191  
 sigwait subroutine 192  
 sigwaitinfo subroutine 191  
 sin subroutine 193  
 sine subroutines  
   sinf 193  
 sinf subroutine 193  
 single-byte conversion 468  
 sinh subroutine 194  
 sinhf subroutine 194  
 sinhl subroutine 194  
 sinl subroutine 193  
 sleep subroutine 196  
 slk\_atroff subroutine 627  
 slk\_init subroutine 630  
 slk\_label subroutine 631  
 slk\_noutrefresh subroutine 631  
 slk\_refresh subroutine 632  
 slk\_restore subroutine 633  
 slk\_set subroutine 633  
 slk\_touch subroutine 634  
 SMIT ACL database 127  
 SNRM2 function 667  
 socketmark subroutine 197  
 soft function key label, restore 633  
 soft function key-label 630  
 soft function key, setup 633  
 soft function key, update 634  
 soft label subroutines 627  
 soft label, label name 631  
 soft label, update 631, 632  
 sputl subroutine 154  
 sqrt subroutine 238  
 sqrtf subroutine 238  
 sqrtl subroutine 238  
 square root subroutines  
   sqrtf 238  
 srand subroutine 5  
 random subroutine 7  
 src error message  
   src error code 241  
 SRC error messages  
   retrieving 240  
 src request headers  
   return address 243  
 SRC requests  
   getting subsystem reply information 241  
   sending replies 250  
 SRC status text  
   returning title line 264  
 SRC status text representations  
   getting 265  
 SRC subroutines  
   src\_err\_msg 240  
   srcrrqs 241  
   srcsbuf 244  
   srcsbuf\_r 247  
   srcsrpy 250  
   srcsrqt 253  
   SRC subroutines (*continued*)  
     srcsrqt\_r 256  
     srcstat 259  
     srcstat\_r 262  
     srcstathdr 264  
     srcstattxt 265  
     srcstattxt\_r 265  
     srcstop 266  
     srcstrt 269  
   src\_err\_msg subroutine 240  
   src\_err\_msg\_r subroutine 241  
   srcrrqs subroutine 241  
   srcrrqs\_r subroutine 243  
   srcsbuf subroutine 244  
   srcsbuf\_r subroutine 247  
   srcsrpy subroutine 250  
   srcsrqt subroutine 253  
   srcsrqt\_r subroutine 256  
   srcstat subroutine 259  
   srcstat\_r subroutine 262  
   srcstathdr subroutine 264  
   srcstattxt subroutine 265  
   srcstattxt\_r subroutine 265  
   srcstop subroutine 266  
   srcstrt subroutine 269  
   SROT subroutine 664  
   SROTG subroutine 663  
   SROTM subroutine 671  
   SROTMG subroutine 672  
   SSBMV subroutine 681  
   SSCAL subroutine 668  
   sscanf subroutine 109  
   ssignal subroutine 271  
   SSPMV subroutine 682  
   SSPR subroutine 701  
   SSPR2 subroutine 704  
   SSWAP subroutine 666  
   SSYMM subroutine 706  
   SSYMV subroutine 680  
   SSYR subroutine 700  
   SSYR2 subroutine 702  
   SSYR2K subroutine 713  
   SSYRK subroutine 710  
   stack, alternate 174  
   standard screen  
     clearing 536  
     refreshing 657  
   standend subroutine 635  
   standout subroutine 635  
   start\_color subroutine 636  
   stat subroutine 277  
   statacl subroutine 272  
   statfs subroutine 274  
   statvfs subroutine 276  
   statx subroutine 277  
   STBMV subroutine 685  
   STBSV subroutine 690  
   store screen coordinates 560  
   STPMV subroutine 687  
   STPSV subroutine 692  
   strcasecmp subroutine 283

- strcat subroutine 281
- strchr subroutine 291
- strcmp subroutine 283
- strcoll subroutine 283
- strcpy subroutine 281
- strcspn subroutine 291
- strdup subroutine 281
- streams
  - assigning buffers 129
- strerror subroutine 285
- strfmon subroutine 286
- strftime subroutine 288
- string conversion
  - strtof 295
  - strtoimax 297
  - strtold 295
  - strtoumax 297
  - to double-precision floating points 514
  - to integers 299, 515
  - to long integers 515
- string manipulation macros
  - varargs 429
- string manipulation subroutines
  - re\_comp 28
  - re\_exec 28
  - strncollen 294
  - wordexp 503
  - wordfree 506
  - wstring 511
- string operations
  - appending strings 281
  - comparing strings 283
  - copying strings 281
  - determining existence of strings 291
  - determining string location 291
  - determining string size 291
  - splitting strings into tokens 291
- string subroutines
  - index 291
  - rindex 291
  - strcasecmp 283
  - strcat 281
  - strchr 291
  - strcmp 283
  - strcoll 283
  - strcpy 281
  - strcspn 291
  - strdup 281
  - strerror 285
  - strlen 291
  - strncasecmp 284
  - strncat 281
  - strncmp 283
  - strncpy 281
  - strpbrk 291
  - strrchr 291
  - strsep 291
  - strspn 291
  - strstr 291
  - strtok 291
  - strtok\_r 298
- string subroutines (*continued*)
  - strxfrm 281
- strings
  - breaking strings into tokens 298
  - compiling for pattern matching 28
  - performing operations on type wchar 511
  - returning number of collation values 294
- strlen subroutine 291
- STRMM subroutine 716
- STRMV subroutine 683
- strncasecmp subroutine 284
- strncat subroutine 281
- strncmp subroutine 283
- strncollen subroutine 294
- strncpy subroutine 281
- strpbrk subroutine 291
- strptime subroutine 301
- strrchr subroutine 291
- strsep subroutine 291
- STRSM subroutine 718
- strspn subroutine 291
- strstr subroutine 291
- STRSV subroutine 688
- strtod subroutine 295
- strtof subroutine 295
- strtoimax subroutine 297
- strtok subroutine 291
- strtok\_r subroutine 298
- strtol subroutine 299
- strtold subroutine 295
- strtoul subroutine 299
- strtoumax subroutine 297
- strxfrm subroutine 281
- stty subroutine 303
- subpad subroutine 637
- subroutines
  - remote statistics interface
    - RSiChangeFeed 53
    - RSiChangeHotFeed 54
    - RSiClose 55
    - RSiCreateStatSet 57
    - RSiDelSetHot 58
    - RSiDelSetStat 59
    - RSiFirstCx 60
    - RSiFirstStat 62
    - RSiGetHotItem 63
    - RSiGetRawValue 65
    - RSiGetValue 66
    - RSiInIt 68
    - RSiInstantiate 69
    - RSiMainLoop 72
    - RSiNextCx 73
    - RSiNextStat 74
    - RSiOpen 75
    - RSiPathAddSetStat 78
    - RSiPathGetCx 79
    - RSiStartFeed 80
    - RSiStartHotFeed 81
    - RSiStatGetPath 82
    - RSiStopHotFeed 85
  - restart behavior 177

- subroutines (*continued*)
  - SPMI interface
    - SpmiAddSetHot 198
    - SpmiCreateHotSet 201
    - SpmiCreateStatSet 202
    - SpmiDdsAddCx 203
    - SpmiDdsDelCx 204
    - SpmiDdsInit 205
    - SpmiDelSetHot 207
    - SpmiDelSetStat 208
    - SpmiExit 210
    - SpmiFirstCx 210
    - SpmiFirstHot 211
    - SpmiFirstStat 213
    - SpmiFirstVals 214
    - SpmiFreeHotSet 215
    - SpmiFreeStatSet 216
    - SpmiGetCx 217
    - SpmiGetHotSet 218
    - SpmiGetStat 219
    - SpmiGetStatSet 221
    - SpmiGetValue 222
    - Spmilnit 223
    - SpmiInstantiate 225
    - SpmiNextCx 226
    - SpmiNextHot 227
    - SpmiNextHotItem 228
    - SpmiNextStat 230
    - SpmiNextVals 232
    - SpmiNextValue 232
    - SpmiPathAddSetStat 234
    - SpmiPathGetCx 236
    - SpmiStatGetPath 237
  - subservers 244, 247
  - substring, wide character 454
  - subsystems
    - getting status 244, 247
    - returning status 259, 262
    - sending requests 253, 256
    - starting 269
    - stopping 266
  - subwin subroutine 638
  - subwindows 637
  - superbox subroutine 531
  - superbox1 subroutine 531
  - supplementary process group IDs
    - setting 134
  - swab subroutine 304
  - swapoff subroutine 305
  - swapon subroutine 306
  - swapping memory
    - activating 305, 306
    - returning information on devices 307
  - swapqpry subroutine 307
  - symbolic links
    - creating 308
    - reading contents 22
  - symlink subroutine 308
  - symmetric operations
    - performing rank 1 700, 701
    - performing rank 2 702, 704
  - symmetric operations (*continued*)
    - performing rank 2k 713
    - performing rank k 710
  - sync subroutine 310
  - synchronize I cache with D cache 311
  - SYS\_CFGDD operation 316
  - SYS\_CFGKMOD operation 318
  - SYS\_GETLPAR\_INFO operation 319
  - SYS\_GETPARMS operation 320
  - SYS\_KLOAD operation 321
  - SYS\_KULOAD operation 323
  - SYS\_QDVSW operation 325
  - SYS\_QUERYLOAD operation 326
  - SYS\_SETPARMS operation 327
  - SYS\_SINGLELOAD operation 328
  - sysconf subroutine 311
  - sysconfig operations
    - SYS\_CFGDD 316
    - SYS\_CFGKMOD 318
    - SYS\_GETLPAR\_INFO 319
    - SYS\_GETPARMS 320
    - SYS\_KLOAD 321
    - SYS\_KULOAD 323
    - SYS\_QDVSW 325
    - SYS\_QUERYLOAD 326
    - SYS\_SETPARMS 327
    - SYS\_SINGLELOAD 328
  - sysconfig subroutine 315
  - syslog subroutine 329
  - syslog\_r subroutine 332
  - system limits
    - determining values 311
  - System Performance Measurement Interface
    - subroutines
      - SpmiAddSetHot 198
      - SpmiCreateHotSet 201
      - SpmiCreateStatSet 202
      - SpmiDdsAddCx 203
      - SpmiDdsDelCx 204
      - SpmiDdsInit 205
      - SpmiDelSetHot 207
      - SpmiDelSetStat 208
      - SpmiExit 210
      - SpmiFirstCx 210
      - SpmiFirstHot 211
      - SpmiFirstStat 213
      - SpmiFirstVals 214
      - SpmiFreeHotSet 215
      - SpmiFreeStatSet 216
      - SpmiGetCx 217
      - SpmiGetHotSet 218
      - SpmiGetStat 219
      - SpmiGetStatSet 221
      - SpmiGetValue 222
      - Spmilnit 223
      - SpmiInstantiate 225
      - SpmiNextCx 226
      - SpmiNextHot 227
      - SpmiNextHotItem 228
      - SpmiNextStat 230
      - SpmiNextVals 232

## System Performance Measurement Interface

(continued)

subroutines (continued)

  SpmiNextValue 232

  SpmiPathAddSetStat 234

  SpmiPathGetCx 236

  SpmiStatGetPath 237

system subroutine 337

## T

t\_rcvreldata

  subroutine 367

t\_rcvv subroutine 369

t\_rcvvudata subroutine 371

t\_sndreldata

  subroutine 376

t\_sndv subroutine 373

t\_sndvudata

  subroutine 377

t\_sysconf subroutine 379

tables

  sorting data 1

tan subroutine 339

tanf subroutine 339

tangent subroutines

  tanf 339

tanh subroutine 340

tanhf subroutine 340

tanh1 subroutine 340

tanl subroutine 339

TCB attributes

  querying or setting 341

tcb subroutine 341

tcdrain subroutine 342

tcflow subroutine 343

tcflush subroutine 344

tcgetattr subroutine 346

tcgetpgrp subroutine 347

tcsendbreak subroutine 348

tcsetattr subroutine 349

tcsetpgrp subroutine 351

tdelete subroutine 410

tempnam subroutine 363

temporary files

  constructing names 363

  creating 363

termcap identifiers

  returning Boolean entry 641

  returning numeric entry 642

  returning string entry 642

termdef subroutine 352

terminal attributes

  getting 346

  setting 349

terminal capabilities

  applying parameters to 643, 651

  insert-character capability 571

  insert-line capability 572

terminal capabilities, disable 557

terminal color support 570

terminal manipulation

  determining number of lines and columns 624, 640

  echoing characters 552

  outputting string with padding information 606, 652

  switching input/output of curses subroutines 624

  toggling new-line and return translation 598

terminal modes

  CBREAK 535

  program 609

  raw 607

  resetting 610

  saving 545

  shell 546, 610

terminal names 413

terminal numeric capability 646

terminal speed 529

terminal string capability 647

terminal states

  getting 303, 346

  setting 303, 349

terminal structures 612

terminal variables 621

terminals

  beeping 530

  delaying output to 549

  determining type 413

  flashing 557

  getting names 413

  putting in video attribute mode 656

  querying characteristics 352

  refreshing 603, 608

  setting up 594

  verbose name 586

terminfo database 644

test\_and\_set subroutine 353

tfind subroutine 410

tgamma subroutine 354

tgammaf subroutine 354

tgammal subroutine 354

tgetent subroutine 640

tgetflag subroutine 641

tgetnum subroutine 642

tgetstr subroutine 642

tgoto subroutine 643

thread\_self subroutine 360

thread\_setsched subroutine 360

Thread-Safe C Library

  subroutines

    rand\_r 6

    readdir\_r 21

    strtok\_r 298

Threads Library

  signal, sleep, and timer handling

    raise subroutine 4

    sithreadmask subroutine 189

  sigqueue subroutine 181

  sigtimedwait subroutine 191

  sigwait subroutine 192

  sigwaitinfo subroutine 191

tigetflag subroutine 644

tigetnum subroutine 646

- tigetstr subroutine 647
- time format conversions 288, 301, 447
- time manipulation subroutines
  - nsleep 196
  - reltimerid 39
  - sleep 196
  - usleep 196
- time stamps
  - trace 400
- time subroutines
  - read\_real\_time 23
  - time\_base\_to\_time 23
- time\_base\_to\_time Subroutine 23
- timeout mode 599
- times subroutine 355
- timezone subroutine 357
- tmpfile subroutine 363
- tmpnam subroutine 363
- touchline subroutine 648
- touchoverlap subroutine 649
- touchwin subroutine 650
- towctrans subroutine 365
- towlower subroutine 366
- towupper subroutine 367
- tparam subroutine 651
- tputs subroutine 652
- trace channels
  - halting data collection 404
  - recording trace event for 400
  - starting data collection 404
- trace data
  - halting collection 404
  - recording 400
  - starting collection 404
- trace events
  - recording 400, 402
- trace sessions
  - starting 405
  - stopping 406
- trace subroutines
  - trcgen 400
  - trcgent 400
  - trchhook 402
  - trchhook64 402
  - trcoff 404
  - trcon 404
  - trcstart 405
  - trcstop 406
  - utrchhook 402
  - utrhook64 402
- transmission of data
  - suspending 343
  - waiting for completion 342
- trc\_close subroutine 380
- trc\_compare subroutine 381
- trc\_find\_first subroutine 381
- trc\_find\_next subroutine 381
- trc\_free subroutine 385
- trc\_hkaddset subroutine 386
- trc\_hkdelset subroutine 386
- trc\_hkemptysubroutine 386
- trc\_hkfillset subroutine 386
- trc\_hkisset subroutine 386
- trc\_hookname subroutine 387
- trc\_libcctl subroutine 388
- trc\_loginfo subroutine 390
- trc\_open subroutine 392
- trc\_perror subroutine 394
- trc\_read subroutine 395
- trc\_seek subroutine 398
- trc\_strerror subroutine 399
- trc\_tell subroutine 398
- trcgen subroutine 400
- trcgent subroutine 400
- trchhook subroutine 402
- trchhook64 subroutine 402
- trcoff subroutine 404
- trcon subroutine 404
- trcstart subroutine 405
- trcstop subroutine 406
- trigonometric functions
  - computing 193
  - computing hyperbolic 194
- trunc subroutine 406
- truncate subroutine 407
- truncf subroutine 406
- truncl subroutine 406
- Trusted Computing Base attributes
  - querying or setting 341
- tsearch subroutine 410
- tty (teletypewriter)
  - flushing driver queue 580
- tty devices
  - determining 413
- tty locking functions
  - controlling 412
- tty modes
  - restoring state 611
  - saving state 614
- tty subroutines
  - setcsmap 131
- ttylock subroutine 412
- ttylocked subroutine 412
- ttyname subroutine 413
- ttyslot subroutine 415
- ttyunlock subroutine 412
- ttywait subroutine 412
- twalk subroutine 410
- type ahead check 653
- type-ahead characters
  - flushing 558
- typeahead subroutine 653

**U**

- ulimit subroutine 415
- umask subroutine 418
- umount subroutine 419
- uname subroutine 420
- unamex subroutine 420
- unctrl subroutine 654
- ungetc subroutine 422

- ungetch subroutine 655
- ungetwc subroutine 422
- unlink subroutine 423
- unload subroutine 424
- unlockpt subroutine 426
- unsigned long integers
  - converting wide-character strings to 463
- update soft labels 631, 632, 634
- uppercase characters
  - converting from lowercase 367
  - converting to lowercase 366
- user database
  - opening and closing 153
- user information
  - getting and setting 426
- usleep subroutine 196
- usrinfo subroutine 426
- ustat subroutine 274
- utime subroutine 428
- utimes subroutine 428
- utmp file
  - finding current user slot in 415
- utrchook subroutine 402
- utrchook64 subroutine 402
- uvmount subroutine 419

## V

- varargs macros 429
- vectors
  - computing constant times vector plus vector 663
  - copying X to Y 665
  - interchanging X and Y 666
  - returning complex dot products 661, 662
  - returning dot products 661, 670
  - returning sum of absolute values 668
  - scaling by constants 668
- VFS (Virtual File System)
  - mounting 435
  - unmounting 419
- vfscanf subroutine 432
- vwprintf subroutine 433
- vwscanf subroutine 433
- vidattr subroutine 656
- video attributes
  - alarm signals
    - beeping 530
    - flashing 557
  - highlight mode 635
  - putting terminal in specified mode 656
  - setting 528
  - turning off 526
  - turning on 528
- vidputs subroutine 656
- Virtual File System 435
- virtual screen cursor coordinates 568
- vmount subroutine 435
- vscanf subroutine 432
- vsprintf subroutine 439
- vsscanf subroutine 432
- vswscanf subroutine 433
- vwscanf subroutine 433
- vwsprintf subroutine 439

## W

- waddch subroutine 523
- waddstr subroutine 524
- wait subroutine 440
- wait3 subroutine 440
- waitid subroutine 443
- waitpid subroutine 440
- watof subroutine 514
- watoi subroutine 515
- watol subroutine 515
- wattroff subroutine 526
- wattron subroutine 528
- wattrset subroutine 528
- wclear subroutine 536
- wclrtoebot subroutine 540
- wclrtoeol subroutine 541
- wcscat subroutine 444
- wcschr subroutine 444
- wcscmp subroutine 444
- wcscoll subroutine 446
- wcscpy subroutine 444
- wcscspn subroutine 444
- wcsftime subroutine 447
- wcsid subroutine 448
- wcslen subroutine 449
- wcsncat subroutine 449
- wcsncmp subroutine 449
- wcsncpy subroutine 449
- wcspbrk subroutine 451
- wcsrchr subroutine 451
- wcsrtombs subroutine 452
- wcsspn subroutine 453
- wcsstr subroutine 454
- wcstod subroutine 455
- wcstof subroutine 455
- wcstoimax subroutine 457
- wcstok subroutine 458
- wcstol subroutine 459
- wcstold subroutine 455
- wcstoll subroutine 459
- wcstombs subroutine 461
- wcstoul subroutine 463
- wcstoumax subroutine 457
- wcswcs subroutine 464
- wcswidth subroutine 465
- wcsxfrm subroutine 466
- wctob subroutine 468
- wctomb subroutine 468
- wctrans subroutine 469
- wctype subroutine 470
- wcwidth subroutine 471
- wdelch subroutine 550
- wdeleteln subroutine 551
- wechochar subroutine 553
- werase subroutine 555
- wgetch subroutine 561
- wgetstr subroutine 566

- wide character format
  - vwscanf 433
  - vswscanf 433
  - vscanf 433
- wide character output 433
- wide character strings
  - wcstof 455
  - wcstoimax 457
  - wcstold 455
  - wcstoumax 457
- wide character subroutines
  - get\_wctype 470
  - towlower 366
  - toupper 367
  - ungetc 422
  - ungetwc 422
  - wcscat 444
  - wcschr 444
  - wcscmp 444
  - wscoll 446
  - wscopy 444
  - wcscspn 445
  - wcsftime 447
  - wcsid 448
  - wcslen 449
  - wcsncat 449
  - wcsncmp 449
  - wcsncpy 449
  - wcspbrk 451
  - wcsrchr 451
  - wcsspn 453
  - wcstod 455
  - wcstok 458
  - wcstol 459
  - wcstoll 459
  - wcstombs 461
  - wcstoul 463
  - wcswcs 464
  - wcswidth 465
  - wcsxfrm 466
  - wctomb 468
  - wctype 470
  - wcwidth 471
- wide character substring 454
- wide character to single-byte 468
- wide character, memory 500, 501, 502, 503
- wide characters
  - comparing strings 446
  - converting
    - from date and time 447
    - lowercase to uppercase 367
    - to double-precision number 455
    - to long integer 459
    - to multibyte 461, 468
    - to tokens 458
    - to unsigned long integer 463
    - uppercase to lowercase 366
  - determining display width 465, 471
  - determining number in string 449
  - locating character sequences 464
  - locating single characters 451
- wide characters (*continued*)
  - obtaining handle for valid property names 470
  - operations on null-terminated strings 445, 449
  - pushing into input stream 422
  - returning charsetID 448
  - returning number in initial string segment 453
  - transforming strings to codes 466
- winch subroutine 574
- window coordinates 560
- window manipulation
  - creating structures
    - pad 592
    - subwindow 638
    - window 596
    - window buffer 587
  - drawing boxes 531
  - marking changed overlap 649
  - overwriting window 601
  - refreshing
    - characters 648, 650
    - current screen 603, 608, 658
    - standard screen 658
    - terminal 603, 608, 658
    - window 649, 657
- window size 565
- window, copy 543
- windows 559
  - clearing 536, 537
  - creating 596, 638
  - deleting 552
  - erasing 555
  - moving 591
  - refreshing 649, 657
  - scrolling 619, 620, 622
  - setting standout bit pattern 551
- winsch subroutine 578
- winsertln subroutine 579
- wmemchr subroutine 500
- wmemcmp subroutine 501
- wmemcpy subroutine 501
- wmemmove subroutine 502
- wmemset subroutine 503
- wmove subroutine 589
- wnoutrefresh subroutine 657
- word expansions
  - performing 503
- wordexp subroutine 503, 506
- wordfree subroutine 506
- wprintw subroutine 604
- wrefresh subroutine 608
- write contents of virtual screen 615
- write operations
  - writing to files 506
- write subroutine
  - described 506
- writex subroutine
  - described 506
- writex subroutine
  - described 506
- wscanw subroutine 614

wsetscrrg subroutine 622  
wsscanf subroutine 109  
wstandend subroutine 635  
wstandout subroutine 635  
wstring subroutines 511  
wstrtod subroutine 514  
wstrtol subroutine 515

ZTRSM subroutine 718  
ZTRSV subroutine 688

## X

xcrypt\_btoa 516  
xcrypt\_decrypt subroutine 516  
xcrypt\_encrypt subroutine 516  
xcrypt\_free subroutine 516  
xcrypt\_hash subroutine 516  
xcrypt\_key\_setup subroutine 516  
xcrypt\_malloc subroutine 516  
xcrypt\_printb subroutine 516  
xcrypt\_randbuff subroutine 516  
XTI variables 379

## Y

yield processor 115  
yield subroutine 520

## Z

ZAXPY subroutine 663  
ZCOPY subroutine 665  
ZDOTC function 661  
ZDOTU function 662  
ZDROT subroutine 664  
ZDSCAL subroutine 668  
ZGBMV subroutine 674  
ZGEMM subroutine 705  
ZGEMV subroutine 673  
ZGERC subroutine 695  
ZGERU subroutine 695  
ZHBMV subroutine 677  
ZHEMM subroutine 708  
ZHEMV subroutine 676  
ZHER subroutine 696  
ZHER2 subroutine 698  
ZHER2K subroutine 714  
ZHERK subroutine 711  
ZHPMV subroutine 679  
ZHPR subroutine 697  
ZHPR2 subroutine 699  
ZROTG subroutine 663  
ZSCAL subroutine 668  
ZSWAP subroutine 666  
ZSYMM subroutine 706  
ZSYR2K subroutine 713  
ZSYRK subroutine 710  
ZTBMV subroutine 685  
ZTBSV subroutine 690  
ZTPMV subroutine 687  
ZTPSV subroutine 692  
ZTRMM subroutine 716  
ZTRMV subroutine 683

## Vos remarques sur ce document / Technical publication remark form

**Titre / Title :** Bull AIX 5L Technical Reference Base Operating System and Extensions Volume 2/2

**N° Référence / Reference N° :** 86 A2 48EF 02

**Daté / Dated :** May 2003

### ERREURS DETECTEES / ERRORS IN PUBLICATION

### AMELIORATIONS SUGGEREES / SUGGESTIONS FOR IMPROVEMENT TO PUBLICATION

Vos remarques et suggestions seront examinées attentivement.

Si vous désirez une réponse écrite, veuillez indiquer ci-après votre adresse postale complète.

Your comments will be promptly investigated by qualified technical personnel and action will be taken as required.

If you require a written reply, please furnish your complete mailing address below.

NOM / NAME : \_\_\_\_\_ Date : \_\_\_\_\_

SOCIETE / COMPANY : \_\_\_\_\_

ADRESSE / ADDRESS : \_\_\_\_\_

Remettez cet imprimé à un responsable BULL ou envoyez-le directement à :

Please give this technical publication remark form to your BULL representative or mail to:

**BULL CEDOC  
357 AVENUE PATTON  
B.P.20845  
49008 ANGERS CEDEX 01  
FRANCE**

# Technical Publications Ordering Form

## Bon de Commande de Documents Techniques

To order additional publications, please fill up a copy of this form and send it via mail to:

Pour commander des documents techniques, remplissez une copie de ce formulaire et envoyez-la à :

**BULL CEDOC**  
**ATTN / Mr. L. CHERUBIN**  
**357 AVENUE PATTON**  
**B.P.20845**  
**49008 ANGERS CEDEX 01**  
**FRANCE**

**Phone / Téléphone :** +33 (0) 2 41 73 63 96  
**FAX / Télécopie :** +33 (0) 2 41 73 60 19  
**E-Mail / Courrier Electronique :** srv.Cedoc@franp.bull.fr

Or visit our web sites at: / Ou visitez nos sites web à:

<http://www.logistics.bull.net/cedoc>

<http://www-frec.bull.com> <http://www.bull.com>

CEDOC Reference # N° Référence CEDOC	Qty Qté	CEDOC Reference # N° Référence CEDOC	Qty Qté	CEDOC Reference # N° Référence CEDOC	Qty Qté
__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]	
__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]	
__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]	
__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]	
__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]	
__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]	
__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]		__ __ __ __ __ [ __ ]	
[ __ ]: <b>no revision number means latest revision</b> / pas de numéro de révision signifie révision la plus récente					

NOM / NAME : \_\_\_\_\_ Date : \_\_\_\_\_

SOCIETE / COMPANY : \_\_\_\_\_

ADRESSE / ADDRESS : \_\_\_\_\_

PHONE / TELEPHONE : \_\_\_\_\_ FAX : \_\_\_\_\_

E-MAIL : \_\_\_\_\_

**For Bull Subsidiaries / Pour les Filiales Bull :**

Identification: \_\_\_\_\_

**For Bull Affiliated Customers / Pour les Clients Affiliés Bull :**

**Customer Code / Code Client :** \_\_\_\_\_

**For Bull Internal Customers / Pour les Clients Internes Bull :**

**Budgetary Section / Section Budgétaire :** \_\_\_\_\_

**For Others / Pour les Autres :**

**Please ask your Bull representative. / Merci de demander à votre contact Bull.**



**BULL CEDOC**  
**357 AVENUE PATTON**  
**B.P.20845**  
**49008 ANGERS CEDEX 01**  
**FRANCE**

**ORDER REFERENCE**  
**86 A2 48EF 02**

PLACE BAR CODE IN LOWER  
LEFT CORNER



Utiliser les marques de découpe pour obtenir les étiquettes.  
Use the cut marks to get the labels.



AIX  
AIX 5L Technical  
Reference  
BIOS and  
Extensions  
Volume 2/2  
86 A2 48EF 02



AIX  
AIX 5L Technical  
Reference  
BIOS and  
Extensions  
Volume 2/2  
86 A2 48EF 02



AIX  
AIX 5L Technical  
Reference  
BIOS and  
Extensions  
Volume 2/2  
86 A2 48EF 02

